The Early-Career Consequences of Digital Technologies

Elliott Weder

November 10, 2026

[Latest version]

How do digital technologies affect early-career employment and wage growth? Using administrative data from the Netherlands for 2006–2024, I document two concurrent trends across school-leaving cohorts: the probability of employment after leaving school fell, but early-career wage growth rose. I propose that the rise of digital technologies offers a unifying explanation for both trends. Following investment spikes in digital technologies, such as software and databases, firms reduce their early-career employment share and increase early-career wage growth. By enabling firms to codify their routines and knowledge, digital technologies substitute for early-career workers but augment the experience of those who do become employed. A calibrated heterogeneous-firm model that embeds both channels shows that the decline in the price of digital technologies is quantitatively important for both cohort trends. Substitution rather than augmentation is the dominant channel.

Keywords: digital technologies; early-career workers; returns to experience; labour demand; heterogeneous firms
JEL codes: J24, J23, O33, J31, E22

Elliott Weder (weder@ese.eur.nl): Erasmus University Rotterdam and Tinbergen Institute. I thank Agnieszka Markiewicz and Eric Bartelsman for their continued support throughout this project. I also thank Edoardo Maria Acabbi, Pierre Cahuc, Pauline Corblet, Aksel Erbahar, Ana Figueiredo, Florian van Genderen, Nicolas Groshenny, Matthias Kehrig, Laura Pilossoph, Anthony Terriau, Tommaso Tulkens, Felix Ward, and participants at AIEL Rome, Erasmus University Rotterdam, ESPE Helsinki, IZA@LISER Summer School, Le Mans, Sciences Po, and the University of Amsterdam for helpful comments. Results are based on calculations using non-public microdata from Statistics Netherlands.

1. Introduction

Early-career outcomes shape lifetime labour-market trajectories. Two-thirds of life-cycle wage growth occurs within the first ten years after leaving school (Topel and Ward 1992), while initial labour-market conditions, including the first employer and the state of the business cycle, can have consequences that persist for decades (von Wachter 2020; Arellano-Bover 2024). Two margins of early-career outcomes are particularly important: whether employment is found and, conditional on employment, how quickly wages grow as experience accumulates. In this paper, I document that these two margins have moved in opposite directions over the past two decades. Finding employment has become more difficult, while the wage returns to early-career experience have increased. I refer to this pattern as the *entry–returns divergence*.

I propose firms' adoption of digital technologies (DT), such as software, databases, and cloud services, as an explanation for this divergence. Technological change has long been recognised as an important driver of changes in employment and wages because it does not affect all jobs and tasks equally (Autor et al. 2003; Acemoglu and Autor 2011; Acemoglu and Restrepo 2018). Over the past two decades, the dominant form of technological change has been the rise of DT (Bontadini et al. 2023). These technologies allow firms to codify their knowledge, data, and routines, potentially substituting for some jobs while complementing others.

I consider two channels through which these technologies may affect early-career outcomes. First, a *substitution* channel: early-career jobs are typically structured and routine (Caicedo et al. 2019; Deming 2025), making them particularly exposed to substitution by DT. As firms adopt these technologies, they may reduce their demand for early-career workers, so that finding employment becomes more difficult. Second, an *augmentation* channel: DT may augment early-career workers' productivity as they gain experience, for example by facilitating skill accumulation or shifting work towards non-routine, judgement-intensive tasks that are rewarded more as experience accumulates (Deming 2021).[1] Either channel could increase the returns to early-career experience, by which I mean wage growth in the first five years after leaving school. Substitution lowers wages at the start of the career, while augmentation raises them later on. Thus, firms' DT adoption may simultaneously reduce early-career employment and increase early-career wage growth.[2]

I begin by documenting the entry–returns divergence using Dutch administrative data covering the universe of individuals' employment histories and educational backgrounds be-

[1] Throughout, *augmentation* refers to factor augmentation: DT directly increase the marginal productivity of experienced labour. Augmentation is a special case of capital–labour complementarity (Krusell et al. 2000).

[2] For a concrete example of the substitution–augmentation asymmetry, consider a law firm that adopts predictive-coding software to search and review documents. The software substitutes for part of what junior associates do, so the firm requires fewer juniors. However, the experience of the juniors who are still hired is augmented, for example because the software helps them accumulate tacit knowledge faster or because their tasks shift towards directing the software and reviewing its output.

tween 2006 and 2024. I define *early career* as the first five years after leaving school. First, across school-leaving cohorts from 2006 to 2019, I find that the early-career employment rate fell by 2.8 percentage points, from 96.3% to 93.5%, after accounting for changes in cohort composition. Second, I find that the returns to early-career experience rose from 65 log points to 74 log points. This increase occurred entirely within firms, rather than through changes in worker mobility or sorting across firms. Both cohort trends were more pronounced among college-educated individuals.

To examine whether firms' adoption of DT contributes to the entry–returns divergence, I link firm-level investment data to individuals' characteristics and labour-market histories. I measure firms' DT investment by their expenditure on automation services, a bookkeeping item that covers software, databases, cloud services, and other automation technologies.[3] I use investment spikes to study how labour demand for early-career workers and the returns to early-career experience change within firms following large, discrete investment increases. My empirical strategy is a stacked difference-in-differences (DiD) event study that compares firms before and after a spike with not-yet-treated firms in the same years.[4] Following an investment spike, the share of early-career workers in firm employment drops from 13.1% to 11.5%. The decline is driven entirely by lower hiring of early-career workers. At the same time, the returns to early-career experience increase by 18% relative to the pre-spike level. Both the employment and the returns effects are more pronounced for college-educated than for non-college-educated individuals. The DiD evidence therefore mirrors the aggregate cohort trends: fewer early-career workers, higher returns to early-career experience, and stronger effects for college-educated individuals.

I develop a quantitative heterogeneous-firm model to measure how much of the aggregate entry–returns divergence the rise of DT can account for, and to separate the substitution and augmentation channels. The model builds on the skill-biased technological change literature (Katz and Murphy 1992; Krusell et al. 2000), extending the framework to distinguish between inexperienced early-career labour and experienced *post-early-career* labour (i.e., labour beyond the first five years after leaving school). DT substitute for early-career labour and augment post-early-career labour. I place this production structure in a dynamic firm problem with irreversible investment to generate the investment spikes observed in the data. The returns to early-career experience are measured by the wage gap between the two labour types, which corresponds to the wage growth associated with becoming experienced. DT raise the returns in two ways: substitution lowers the marginal product and wage of early-career labour, whereas augmentation raises the marginal product and wage of post-early-career labour. I calibrate the

[3] The item excludes capitalised investment in machinery, such as industrial robots.

[4] The investment-spike design was pioneered by Bessen et al. (2025), who study displacement and earnings effects on incumbent workers in the same Dutch data that I use. The design is also used in Aghion et al. (2025), Bijnens et al. (2026), and others. An investment spike occurs when a firm's investment exceeds three times its own average, as in Bessen et al. (2025). The average spike is a fourfold increase in investment.

model to match the DiD evidence.

I then use the calibrated model to study the aggregate effects of the rise of DT over the last two decades. When I feed in the observed decline in the relative price of DT investment, firms respond by investing more, and the aggregate DT stock rises. As a result, early-career employment falls by 4.5% and the returns to early-career experience rise by 3.0 log points. Through the lens of the model, the lower DT price more than accounts for the decline in early-career employment across school-leaving cohorts (158%) and explains 33% of the increase in the returns to early-career experience.[5] On both margins, substitution rather than augmentation is the dominant channel.

This paper makes two main contributions. First, I document a novel empirical pattern, the entry–returns divergence: employment declined across school-leaving cohorts while the returns to early-career experience increased. I show that this divergence is particularly pronounced among college-educated individuals and that the increase in the returns to early-career experience occurred entirely within firms, rather than through worker mobility or sorting across firms, which the literature has identified as central to early-career wage growth (Topel and Ward 1992; Bagger et al. 2014).

Second, I provide an explanation for this divergence based on firms' adoption of DT. I present firm-level evidence that DT investment reduces hiring of early-career workers and increases the returns to early-career experience within firms. A quantitative model shows that rising investment in these technologies can account for a substantial share of the aggregate divergence, with substitution away from early-career labour dominating augmentation of experienced labour.

Related literature. This paper contributes to two strands of literature. First, it relates to the literature on technological change and labour demand. A large body of work documents how information technologies have shifted labour demand across skill and occupation groups (Autor et al. 1998, 2003; Goldin and Katz 2008; Acemoglu and Autor 2011; Autor and Dorn 2013; Goos et al. 2014; Akerman et al. 2015), while more recent studies examine the displacement effects of industrial robots (Acemoglu and Restrepo 2020) and artificial intelligence (Brynjolfsson et al. 2025; Humlum and Vestergaard 2025; Althoff and Reichardt 2026; Hosseini and Lichtinger 2026). I contribute to this literature by showing that labour-market experience is another important dimension along which technological change affects the labour market. In particular, DT can reduce employment opportunities for inexperienced individuals while increasing the returns to early-career experience, including among college graduates. These findings complement the literature on skill-biased technological change (Katz and Murphy 1992; Krusell et al. 2000) by showing that technological change can also be experience-biased: even

[5] Forces outside the model, such as the growth of flexible work arrangements and the shift towards services, may cushion the decline in the data.

within education groups, technology can simultaneously substitute for early-career workers but increase the returns to early-career experience conditional on employment.

On the empirical front, the most closely related studies are Beaudry et al. (2016), who document declining demand for cognitive tasks and the resulting occupational downgrading of young college graduates, and Barth et al. (2023), who examine how software investment affects the age–earnings profile. Relative to these studies, I show that DT adoption simultaneously reduces early-career employment opportunities, especially for college-educated entrants, and increases the returns to early-career experience.

On the quantitative front, I build on the skill-biased technological change literature (Katz and Murphy 1992; Krusell et al. 2000) and on the capital–experience complementarity in Jaimovich et al. (2013). I depart from these frameworks in two important ways. First, I incorporate endogenous and lumpy investment in DT. This allows the model to generate the investment spikes examined in the DiD analysis and the aggregate rise in DT investment through firms' investment decisions. Second, DT operate through two separate channels, substitution of early-career labour and augmentation of post-early-career labour, capturing how technological change affects early-career employment and the returns to early-career experience.

This paper also contributes to the literature on early-career outcomes (Topel and Ward 1992; Kahn 2010; Oreopoulos et al. 2012; Altonji et al. 2016; Chetty et al. 2017; Ashworth et al. 2021; Rothstein 2023; Corblet 2024; Corinth and Larrimore 2024; Argan et al. 2025; Deming 2025; Bianchi and Paradisi 2026) in four ways. First, I compare the *within-person* returns to early-career experience across school-leaving cohorts, which allows me to flexibly control for unobserved heterogeneity, composition, and selection. Second, I build on Rothstein (2023), who documents a divergence between employment and wage levels for recent cohorts of young college-educated Americans using US survey data. I show that a related divergence appears in Dutch administrative data, is especially pronounced among college-educated individuals, and has a firm-level mechanism. Third, by combining matched employer–employee data with register data on individuals *outside* the labour market, I show that this divergence operates through the extensive margin, complementing the evidence on age–wage gaps in Bianchi and Paradisi (2026). Fourth, whereas Jeong et al. (2015) attribute changes in the returns to experience over 1968–2007 in the US primarily to supply factors, I find that demand-side factors account for a substantial share of the rise in the returns to early-career experience in the Netherlands from 2006 onwards.

Roadmap. Section 2 documents the entry–returns divergence in the data. Sections 3 and 4 use investment spikes to study firm- and worker-level outcomes in a DiD framework. Section 5 builds and calibrates a heterogeneous-firm model to quantify the contribution of the rise of DT to the entry–returns divergence. Section 6 concludes.

2. Motivating Early-Career Facts

This section documents two long-run changes in early-career outcomes across school-leaving cohorts. The probability of entering employment after schooling fell between the 2006 and 2019 cohorts (Section 2.2), while the returns to early-career experience rose (Section 2.3). Together these two facts form the entry–returns divergence. The remainder of the paper then shows that the rise of DT offers a unifying explanation for the divergence.

2.1. Institutional Background and Individual Data

I begin by describing the Dutch institutional setting and the data sources that I use in my analysis. The data are particularly suited to my research question because I observe (i) the universe of Dutch residents annually, so I can follow individuals' early-career trajectories, (ii) individuals' education level and field of study, which allows me to measure exposure to DT, (iii) the link between workers and their employers, and (iv) firms' investments in DT in a panel data setting, so I can follow firms' investments over time.

Education. Students in the Dutch education system are sorted into academic tracks at the beginning of secondary school based on academic ability. A student's track determines their access to vocational training, universities of applied sciences, or research universities. In the data, I classify individuals whose highest qualification is from an applied or research university as *college*-educated and the remainder as *non-college*-educated.

Labour market. Flexible and part-time work are common in the Netherlands, especially among young workers (Eurostat 2026a,b). Between 70 and 80% of employees are covered by collective bargaining agreements (CBAs), which are often extended to entire sectors. Typically, CBAs specify minimum *wage increases* rather than wage levels, which leaves firms room to set their own wages. The national statutory minimum wage rises with age. Before 2017, the full adult rate was reached at age 23. The age threshold fell from 23 to 21 between 2017 and 2019.

Individual data. I link tax, education, and demographic registers from Statistics Netherlands through unique individual identifiers. The registers cover every resident annually from 2006 to 2024.[6] The tax register reports employment, self-employment, and benefit income, together with an individual's main socio-economic status (employee, self-employed, student, etc.). The

[6] I drop individuals who emigrated, immigrated, or died during the sample period to ensure I observe individuals' entire employment or non-employment history in the Netherlands. Furthermore, I drop individuals who are ever self-employed (fewer than 0.5% of early-career individuals) and anyone ever in the top 0.5% of the pooled income or wage distribution.

education register records each individual's highest qualification and the corresponding field of study. I also observe age, gender, and migration background.

My analysis follows cohorts through their first years after schooling. An individual i belongs to cohort c if c is the final calendar year in which they are a student.[7] I then follow each individual annually for five years, from their cohort year $t = c$ until $t - c = 5$. Potential experience is measured as $\tau \equiv t - c$. Since the data run from 2006 to 2024, I choose $\tau \leq 5$ to balance the number of cohorts with the number of years that I observe individuals after schooling. The final sample spans fourteen cohorts from 2006 to 2019, observed from 2006 until 2024, and six observations per individual ($\tau \in \{0, 1, \ldots, 5\}$). Results are robust to following individuals up to $\tau = 7$ and $\tau = 10$, with fewer cohorts accordingly.

Matched employer–employee data record earnings and hours worked if an individual is employed.[8] The main outcome variable is the log hourly wage, defined as total annual earnings over total annual hours worked. All nominal variables are deflated to 2015 euros. Furthermore, I retain each individual's primary employer per year, defined as the employer paying the largest share of earnings if there are multiple employers, and characteristics of the employment contract at that employer.

I construct a similar panel of individuals whose cohort c is the calendar year in which they are 45 years old. I observe each individual in the age-45 group from 45 years old until 5 years later. The age-45 group acts as a comparison group for the early-career panel because their outcomes are unaffected by early-career dynamics and only reflect aggregate wage growth. I find similar results with an age-50 and age-55 comparison group. The final sample comprises 1.8m early-career and 2.3m age-45 individuals, for a total of 24.6m individual-year observations. Summary statistics are shown in Appendix A.1.

2.2. Fact 1: Decline in Labour-Market Entry Rates

Here I document that the average employment probability in the five years after schooling fell across each successive cohort. To do so, I adapt the approach in Rothstein (2023) to estimate cohort employment effects.

Empirical design. I compare the average employment rate across different cohorts at the same age, with the same education, in the same year. The baseline regression is given by:

$$\tilde{E}_{it} = \sum_{b \neq b_0} \beta_b I_b + \upsilon_{ae} + \lambda_t + \zeta_g + \nu_m + u_{it}. \tag{1}$$

[7] The tax register classifies an individual as a student if they are enrolled in an educational institution and earn below roughly €14.5k (2015 euros). The ceiling excludes nominally enrolled individuals who have already entered the labour market. An individual is not required to have completed a qualification in this final year.

[8] Earnings are gross wages, social-insurance contributions, bonuses, and in-kind payments. Hours are contractual plus overtime hours.

The variable I_b indicates that individual i belongs to cohort bin b. For parsimony, I bin adjacent cohorts (2008–09, 2010–11, etc.), so that a coefficient β_b measures the difference in the average employment rate of cohort bin b relative to the 2006–07 reference bin b_0, adjusted for compositional changes.[9] The terms υ_{ae}, λ_t, ζ_g, and ν_m stand for education (level and field)-by-age, calendar year, gender, and migrant-background fixed effects.[10]

The outcome variable equals $\tilde{E}_{it} = E_{it} - \hat{g}(e, \tau)$. The indicator E_{it} equals one if i is employed at any point in year t, and zero otherwise. The part $\hat{g}(e, \tau)$ captures an employment *ramp* that occurs along τ. Employment is lower in the year someone leaves school than in the years after, because many people do not start work straight away (Appendix B.3). Later cohorts also leave school at older ages (Appendix A.1). With the inclusion of age fixed effects, later cohorts are observed at lower τ, and so lower on the ramp for a given age. The regression attributes this lower employment to the cohort effects β_b. To circumvent this problem, but still control for life-cycle effects, I estimate the ramp $\hat{g}(e, \tau)$ directly and subtract it from the outcome in the main regression. Appendix B.1 contains more details.

I estimate Equation (1) separately for the full early-career sample and for several subsamples: college- and non-college-educated, and DT-exposed and non-exposed fields of study (details below and in Appendix A.4), and within and beyond two years of leaving school.

Each coefficient β_b is identified by comparing individuals with the same qualification and age, say mathematics master's graduates aged 25, across school-leaving cohorts. The education-by-age effects υ_{ae} control for life-cycle effects, such as marriage, parenthood, or home ownership, that can also differ by education level and field of study. These effects also control for differences in educational attainment across cohorts (Table A1). The calendar year effects λ_t absorb transitory macroeconomic conditions common to all early-career individuals, including the Global Financial Crisis and the COVID-19 pandemic.[11] Raw employment profiles by cohort are shown in Appendix B.3. For reference, the average employment rate over $\tau \in \{0, \ldots, 5\}$ for the reference group b_0 is 96.3%.

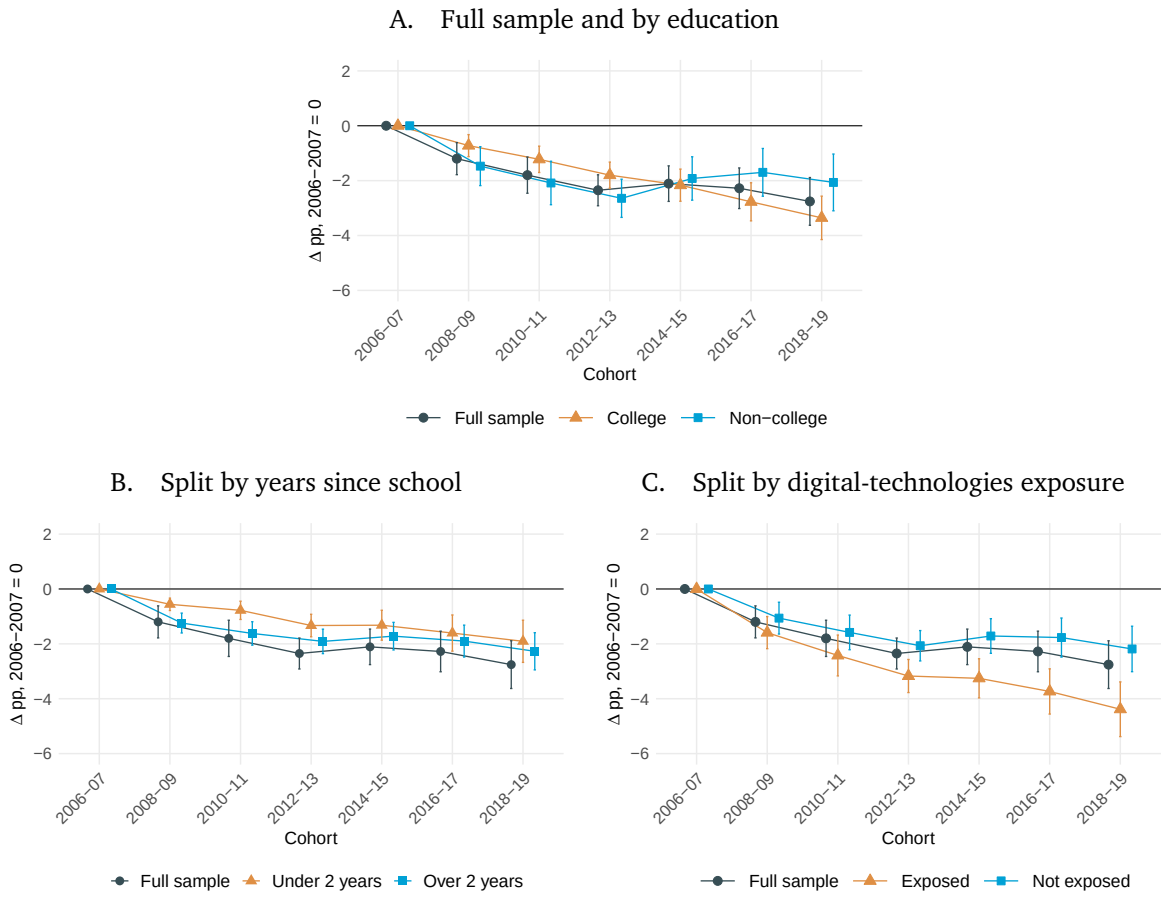
Results. Figure 1 plots the cohort employment effects $\hat{\beta}_b$. In Panel A, the 'Full sample' line shows the overall employment probability fell steadily across early-career cohorts. An individual from the 2018–19 cohorts is 2.8 pp less likely to be employed in their first five years than a comparable individual from the 2006–07 cohorts. The estimates in Panel A are broadly consistent with Rothstein (2023) for young American college graduates.

[9] Estimates for single-year cohorts c are similar and shown in Appendix B.2.

[10] Field of study is recorded for every qualification. Vocational qualifications carry a field (e.g. engineering, health, business) and secondary-school qualifications fall under 'Generic programmes and qualifications'. Education level refers to 'High school or below', 'Vocational degree', 'Bachelor's degree', and 'Master's degree'.

[11] Borjas (1999), Bratsberg et al. (2006), and Rothstein (2023) note that cyclical effects may differ across groups. Since my estimation sample in this subsection contains only early-career individuals, the scope for differential time effects is limited.

FIGURE 1. Cohort effects on early-career employment probabilities.



Notes: The figure plots the estimated cohort employment effects $\hat{\beta}_b$ from Equation (1), estimated by ordinary least squares. The outcome is an indicator for being employed at any point in the year, adjusted for the qualification-specific employment ramp over the years since school (Section 2.2), so estimates are in percentage points relative to the 2006–07 reference cohorts. Adjacent cohorts are binned (2008–09, 2010–11, and so on). Panel A reports estimates for the full sample and separately by education level. Panel B splits the sample into individuals within two years after schooling ($\tau \leq 2$) and beyond ($\tau \geq 3$). Panel C splits by field-of-study exposure to digital technologies (mapping shown in Appendix A.4). The sample is the early-career panel described in Section 2.1: the 2006–2019 cohorts observed for $\tau \in \{0, \ldots, 5\}$. Bars show 95% confidence intervals with standard errors clustered at the cohort level, meaning there are fourteen clusters. Wild cluster bootstrap confidence intervals are shown in Appendix B.6. Number of observations: 10.6 million (full sample), 4.5 million (college), 6.1 million (non-college), 5.3 million (each years-since-schooling split), 2.7 million (exposed), and 7.9 million (non-exposed).

Each panel of Figure 1 also plots estimates $\hat{\beta}_b$ for different subsamples. Panel A highlights that the decline in the early-career employment probability was more pronounced for college graduates than for non-college graduates by 2018–19. Panel B splits the estimates by years since schooling (τ), showing a similar decline at both time horizons. Lastly, Panel C splits the estimates by exposure to digital technologies based on field of study. Exposed fields include computer science, mathematics and natural sciences, and law and business services. Non-exposed fields include healthcare, education, and other similar fields.[12] Panel C shows the cohort decline in employment probability was larger for exposed fields than non-exposed fields.

Robustness. Three checks confirm that the decline is a cohort effect rather than an artefact of the specification or selection. First, the employment rate of the age-45 comparison group did not fall over the same cohorts. Appendix B.4 shows the employment rate of the age-45 group increased by 2.1 pp across cohorts, so the result is specific to labour-market entrants rather than a secular fall in employment.[13] Second, bounding the cohort estimates for selection on unobservables following Oster (2019) leaves the decline intact (Appendix B.7). Third, changing the fixed effects or assumptions on the business cycle effects does not change the results (Appendix B.8). A final concern is that individuals time their labour-market entry to enter in a favourable macroeconomy. However, this would bias the OLS estimates upwards and understate the employment rate decline.

2.3. Fact 2: Increase in the Returns to Early-Career Experience

Next, I document that the wage returns to early-career experience rose across the same cohorts. To that end, I regress the log hourly wage on potential experience τ, controlling for business cycle and compositional changes across cohorts, building up my preferred specification in three steps.

Empirical design. Consider first a regression of log hourly wages $\ln w_{it}$ on potential experience indicators I_τ, similar to Deming (2025):

$$\ln w_{it} = \alpha_i + \sum_{\tau=1}^{5} \beta_\tau I_\tau + u_{it}, \tag{2}$$

in which α_i represent individual fixed effects. For each $\tau \in \{1, \ldots, 5\}$, β_τ measures the cumulative wage return to τ years of potential experience, measured in log points.[14]

[12] The full mapping is described in Appendix A.4. The heterogeneity across exposed and non-exposed fields is robust to changing the mapping.

[13] Moreover, I find the opposite-signed experience and education gradients over age-45 cohorts compared to early-career cohorts (see Appendix B.4).

[14] Using actual experience gives similar results (Appendix B.14).

Second, I interact I_τ with cohort-group indicators I_b to measure the change across cohorts:

$$\ln w_{it} = \alpha_i + \sum_{\tau=1}^{5} \beta_\tau I_\tau + \sum_{b \neq b_0} \sum_{\tau=1}^{5} \rho_{b\tau} (I_\tau \times I_b) + u_{it}. \tag{3}$$

The coefficients β_τ measure the returns to experience for the reference cohort group b_0, whereas $\rho_{b\tau}$ capture the difference in the returns at τ between cohort groups b and b_0.

The coefficient estimates in Equation (3) potentially confound business cycle conditions with secular changes in the returns to early-career experience across cohorts. Year effects λ_t are collinear with the potential experience dummies I_τ once individual fixed effects are included. Therefore, in the third step, I net out the cyclical wage effects using the wage paths of the age-45 individuals. Letting I_y denote the early-career group, I estimate:

$$\ln w_{it} = \alpha_i + \phi_{c\tau} + \sum_{\tau=1}^{5} \eta_\tau (I_\tau \times I_y) + \sum_{b \neq b_0} \sum_{\tau=1}^{5} \theta_{b\tau} (I_\tau \times I_b \times I_y) + u_{it}, \tag{4}$$

in which $\phi_{c\tau}$ are cohort-by-experience effects.[15] Under the assumption that the wage profile of age-45 workers reflects aggregate wage growth and not the returns to early-career experience, differencing against them nets out the cyclical component of the returns to early-career experience.[16] In Equation (4), the coefficients η_τ measure the returns to early-career experience for the reference cohorts b_0. The coefficients $\theta_{b\tau}$ capture the change for group b relative to b_0.[17] For parsimony, I split the fourteen cohorts in half, $b_0 = \{2006, \ldots, 2012\}$ and $b = \{2013, \ldots, 2019\}$. Results for adjacent cohorts are in Appendix B.10. I cluster the standard errors at the cohort level. Raw wage growth profiles by cohort are shown in Appendix B.9.

Results. Panel A of Figure 2 plots the estimates $\hat{\eta}_\tau$, which represent the returns to early-career experience for the 2006–12 cohorts. For that group, hourly wages grew by around 65 log points in the five years after schooling. The magnitude and concave shape of the profile are broadly comparable to what is documented for the US in Deming (2025).

Panel B plots the estimates $\hat{\theta}_{b\tau}$. The 'Full sample' estimates show that the returns to early-career experience steepened across cohorts: the return to the first year of experience rose by 5.6 log points (5.8%) and the cumulative five-year return rose by 9.3 log points (9.7%).[18] Together, the decline in early-career employment rates and the increase in the returns to early-career

[15] I use c here rather than the groups b to net out the business cycle growth as strictly as possible.

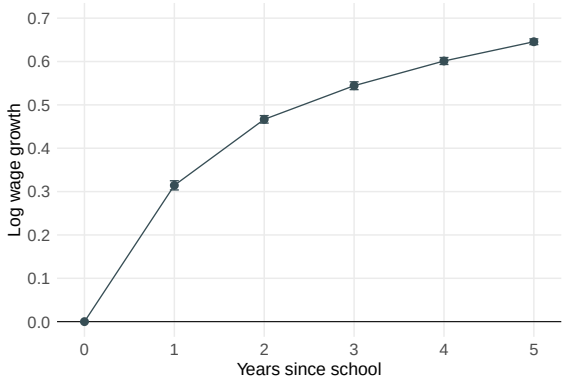
[16] The results are unchanged when I interact the log aggregate unemployment rate with I_y in Appendix B.15 (Borjas 1999; Bratsberg et al. 2006; Rothstein 2023), and when age-50 or age-55 workers form the control group.

[17] If the rise of DT or other mechanisms increases the wage growth of age-45 workers, then the estimates for $\theta_{b\tau}$ are biased towards zero, meaning that I estimate a lower bound.

[18] The $\tau = 1$ return potentially combines the first hire and first year on the job. Both channels represent genuine returns to *labour-market experience*. Nevertheless, below I document that the steepening reflects wage growth within firms rather than changes in hiring wages (at $\tau = 1$) or match composition.

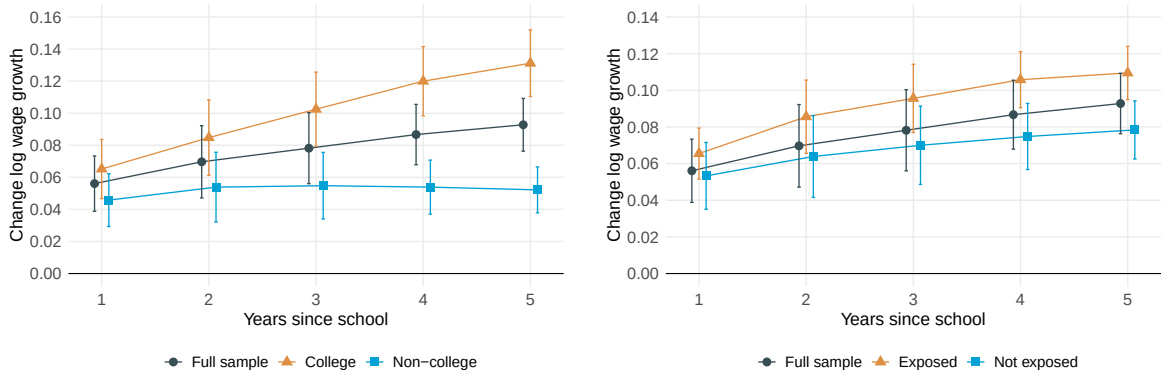
FIGURE 2. Cohort wage returns to early-career experience.

A. Returns for 2006–12 reference cohorts



B. Change by education, 2013–19 cohorts

C. Change by exposure, 2013–19 cohorts



Notes: Panel A plots the estimated returns to early-career experience $\hat{\eta}_\tau$ for the 2006–12 reference cohorts b_0 from Equation (4), estimated by ordinary least squares. Panels B and C plot the estimates $\hat{\theta}_{b\tau}$ for the change in the returns for the 2013–19 cohorts relative to the reference cohorts. I plot the full-sample and separate estimates by education level (Panel B) and by field-of-study exposure to digital technologies (Panel C; mapping in Appendix A.4). The outcome is the log hourly wage, so estimates are in log points. All estimates net out the wage growth of the age-45 comparison group over the same calendar years. The sample is the early-career and age-45 panels described in Section 2.1, restricted to employed individual-years. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Wild cluster bootstrap standard errors are similar (Appendix B.11). Education coverage is not complete for the entire age-45 panel. Number of observations: 21.4 million (Panel A and full sample), 6.9 million (college), 10.4 million (non-college), 4.3 million (exposed), and 13.0 million (non-exposed).

experience constitute the entry–returns divergence.

A Wald test rejects the null hypothesis that all full-sample estimates $\hat{\theta}_{b\tau}$ are equal (p = 0.000): the increase in returns to early-career experience is concave (Panel B). Appendix B.12 further demonstrates that the increase in the returns to early-career experience persisted until $\tau = 7$ and $\tau = 10$, with the caveat that I compare fewer cohorts over time.

Panels B and C show heterogeneity by education and DT-exposure. Mirroring the employment results, the steepening was more pronounced for college graduates and for DT-exposed

fields of study. In contrast to the college-educated profile in Panel B, the non-college profile is flat after the shift up at $\tau = 1$ ($p = 0.601$ for joint equality at all $\tau \geq 1$). Moreover, the increase was more pronounced for DT-exposed fields of study compared to non-exposed fields at each τ ($p < 0.001$).

Within-firm steepening, not reallocation. The steepening occurred within firms rather than through recent entrants sorting into higher-growth firms. In Appendix B.18, I map residualised wage growth into the decomposition framework of Foster et al. (2001). The decomposition enables me to separate sorting and within-firm changes in accounting for the rise in the returns to experience. I find that over 100% of the rise in returns at every τ is accounted for by steepening within firms. Therefore, reallocation across firms offsets part of the increase. Had recent cohorts increasingly matched into high-growth firms through better screening or search, the gain would have loaded onto reallocation. The rise in returns instead reflects a shift in how employers reward experience.[19]

Robustness. The coefficients $\theta_{b\tau}$ in the preferred specification in Equation (4) are identified using only wage growth, not wage levels, because individual fixed effects are included. Therefore, changing selection into employment based on wage levels across cohorts is absorbed by the fixed effects, leaving the estimates unaffected. Appendix B.18 adapts the Abowd et al. (1999) model to separate individual-level and firm-level wage growth. According to the exercise, selection on individual-level wage growth in recent cohorts is not able to explain the results. Furthermore, I bound the change in selection on wages by imputing the wages of the non-employed in Appendix B.17, following Horowitz and Manski (2000). The worst-case imputation bound for the change in the returns to early-career experience is positive and significant. All other robustness checks are collected in Appendix B.

2.4. Digital Technologies and Early-Career Outcomes

The previous two subsections show that labour-market entry and the returns to early-career experience diverged between the 2006 and 2019 cohorts. I propose that the rise of DT can explain the entry–returns divergence. Workers typically begin their careers in structured and routine jobs. Deming (2025) demonstrates that the routineness of occupations declines with labour-market experience. Similar evidence shows that senior consultants perform more customised work than junior consultants (Brandon-Jones et al. 2016; Alvehus 2017). DT allow firms to codify their routines and knowledge, thereby substituting for early-career workers performing codifiable tasks and reducing labour-market entry. At the same time, DT may

[19] The decomposition result is consistent with, e.g., Baksy et al. (2026), who show that job-to-job mobility, and the wage growth associated with these moves, has declined in the US over time.

augment early-career workers' productivity as they gain experience, for example by facilitating skill accumulation or shifting work towards non-routine, judgement-intensive tasks that are rewarded more as experience accumulates (Deming 2021). Both channels plausibly steepen the wage returns to early-career experience.

Selection, supply, and mean reversion. Selection is unlikely to explain my results. First, selection on wage levels is absorbed by the individual fixed effects. Second, worst-case-bound imputation for non-employed wages leaves the rise intact at every τ (Appendices B.17). Third, I show that worker composition accounts for at most one third of the steeper growth in a modified Abowd et al. (1999) decomposition (Appendix B.18).

Changes in supply are also unlikely to explain my results. A labour-supply explanation is possible in principle, because a contraction in the supply of early-career workers would lower labour-market entry. However, in a standard Katz and Murphy (1992) relative-supply model, that increase in relative scarcity would simultaneously raise early-career wages and flatten the wage return profile. If the over-supply of early-career workers pushes down their wages, supply still does not explain why wage growth quickens within firms. Moreover, mean reversion from lower starting wages at $\tau = 0$ cannot explain all of the increase in the returns to early-career experience (Appendix B.16).

Dutch labour-market institutions. Dutch labour-market institutions do not account for the divergence either (Appendix B.19). Within-sector variation accounts for 87–93% of the variance in the increase in the returns to early-career experience, which rules out sector-level CBAs as the driver, since they typically apply to the entire sector. Furthermore, I find that the returns to early-career experience rose more for workers with permanent contracts than with temporary contracts, so my results are not driven by wage gains at conversion from a temporary to a permanent contract (Bolhaar et al. 2018; Daruich et al. 2023; Figueiredo et al. 2026).

In the next two sections, I link firm-level DT investment data to workers' characteristics and labour-market histories and show that DT produce an entry–returns divergence *within firms*, mirroring the aggregate cohort trends. Finally, Section 5 calibrates a quantitative model to the DiD estimates to quantify the total contribution of the rise of DT to the cohort trends and to quantify the relative importance of substitution and augmentation.

3. Digital Technologies and Firms' Demand for Early-Career Workers

In this section, I study how firms change their demand for early-career workers around large spikes in DT investment, the first side of the entry–returns divergence. Section 3.1 introduces

the data that I use, Section 3.2 discusses the empirical strategy, Section 3.3 considers threats to identification, and Section 3.4 presents the results.

3.1. Firm Investment in Digital Technologies

I use the Production Statistics (PS) panel survey to measure firms' DT investment. The PS runs from 2000 until 2023 and is a mandatory survey for firms with at least 50 employees, whereas smaller firms are randomly sampled. Surveyed firms cover construction, manufacturing, wholesale, retail, transportation, and services sectors and employ around 50% of early-career workers in a given year. The survey records various items, such as revenues, costs, and investments. Unique identifiers link firms across years and to workers via the employer–employee data.

I measure DT investment by firms' expenditures on third-party automation services, such as software, databases, and cloud services.[20] These purchases help firms to automate production and their internal procedures.[21] While the item cannot be further disaggregated, Bessen et al. (2025) and Bartelsman et al. (2026) show that expenditure is related to spending on enterprise software and licensing (ERP, CRM, data infrastructure), the use of big data, cloud services, AI, and process innovation more generally. Firms across all sectors invest in DT (Bessen et al. 2025).

The variable specifically excludes capitalised tangible automation technologies, such as industrial robots and computers, and other forms of intangible assets, like brand capital, patents, and R&D. Bessen et al. (2025) find that DT expenditures exceed the value of industrial robot imports.

The advantage of using the survey item to measure DT is threefold. First, the panel structure allows me to track firms' investments over time and, therefore, to identify the timing of investment events for the DiD design. Second, I can link the investment data directly to workers through the administrative data's unique firm and worker identifiers. Third, I do not need to infer the market value of the DT investment because the expenditure item is an official bookkeeping entry for firms (McGrattan and Prescott 2010; Bessen et al. 2025).

Appendix A.2 reports summary statistics for the firm panel. One central feature of the data is that DT investment has risen sharply in recent decades: average investment per worker grew by 70% over the sample, from €1,200 in 2000 to roughly €2,000 in 2023 (Appendix A.3).

[20] Firms in the Netherlands rely mainly on third-party suppliers for software development and maintenance, rather than in-house development (Statistics Netherlands 2016). For example, in 2015 50–70% (depending on the sector and the category) of surveyed firms responded that they relied on external suppliers rather than in-house development for maintenance of ICT infrastructure, design and development of business software, software support, web application support, and data protection. Thus, the majority of firms' DT expenditure is captured by the survey item, despite the fact that it includes third-party spending only.

[21] Statistics Netherlands names the item "automation costs".

3.2. Empirical Strategy

The main outcome of interest is firm j's early-career employment share in year t (for brevity *EC employment share* or *EC share*):

$$y_{jt} = \frac{L_{Ejt}}{L_{Ejt} + L_{Rjt}}.$$

L_{Ejt} denotes the number of early-career workers within five years of potential experience ($\tau \in \{0, \ldots, 5\}$) in the sample used in Sections 2.2 and 2.3. L_{Rjt} represents the remaining workers in firm j subject to the same sample restrictions as in Section 2.1.[22]

The ideal experiment would exogenously shift firms' investment and compare relative labour demand for early-career workers afterwards. An OLS regression of the EC share on investment will not approximate the experiment because investment is lumpy: 31% of firm-years report no investment (Appendix A.2). A firm with zero investment may therefore be a firm that never invests or a firm that made a large outlay the year before.

Instead, I approximate the ideal experiment by exploiting the lumpiness of investment. Following Aghion et al. (2025), Bessen et al. (2025), and Bijnens et al. (2026), I employ a stacked DiD event-study design around large investment spikes to identify how the EC share adjusts. The idea is that spikes (defined below) capture sharp increases in investment, so comparing firms that spike to similar firms that do not will isolate the effect of DT on firms' demand for early-career workers. Observed spikes in investment are consistent with models of non-convex adjustment costs (Cooper and Haltiwanger 2006; Bessen et al. 2025; Bijnens et al. 2026).

Following Bessen et al. (2025), firm j has a spike, or *event*, in year s when its investment I_{js} relative to average total costs $\bar{C}_j$ exceeds three times its own leave-one-out mean:

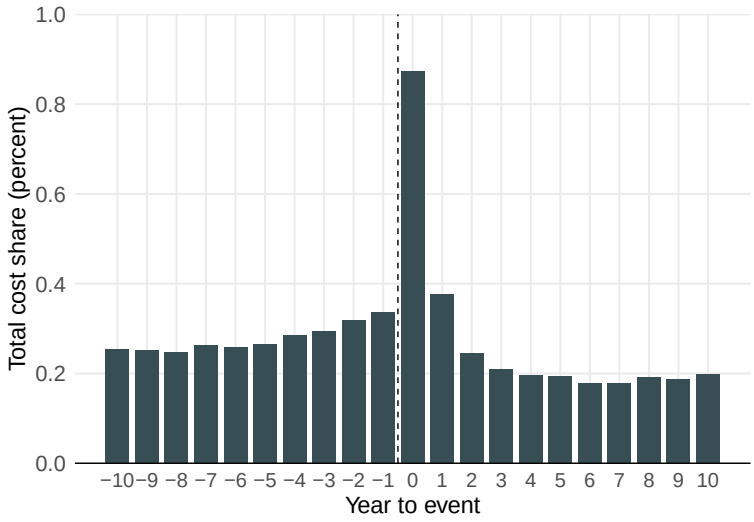
$$I_{\text{spike}_{js}} = \mathbb{1}\left(\frac{I_{js}}{\bar{C}_j} \geq 3 \times \frac{1}{T_j - 1} \sum_{t \neq s} \frac{I_{jt}}{\bar{C}_j}\right), \tag{5}$$

where T_j is the number of years firm j is in the panel. The leave-one-out threshold is firm-specific, so a spike is large relative to the firm's own history rather than to an economy-wide benchmark. I use *spike* and *event* interchangeably hereafter. If a firm experiences multiple spikes, I retain the first one.

Figure 3 shows that investment spikes are discrete. The DT investment to total cost share is mostly flat leading up to a spike, then nearly quadruples in the spike year, before immediately settling back to the pre-spike level. Average investment is around €200k higher in the spike year than in the year prior.

[22] L_{Rjt} excludes workers still in school. I require firms to be observed at least three times in the panel data to appear in my sample.

FIGURE 3. Investment share of total operating costs in event time.



Notes: The figure plots the mean share of digital-technologies investment in total operating costs by event time $k = t - s$ around firms' first investment spike. The sample is the treated firms in the stacked panel described in Section 3.2, with event-years s from 2010 to 2019. Number of firms: 5,334.

The main specification is a stacked DiD that compares spiking firms to firms that have not yet spiked. For each event-year s in which firms can have a spike, I construct a window in calendar time: $t \in \{s-4, \ldots, s+3\}$, with event time $k \equiv t - s$ and $k \in \{-4, \ldots, 3\}$. Treated firms ($I_T = 1$) have their first event in year s, whereas control firms ($I_T = 0$) are *not-yet-treated*, meaning their first event occurs after the window at $t > s + 3$. Both treated and control firms are observed in the same calendar years in the window. I stack each event-year s's panel into a single panel. The final panel, with $s \in \{2010, \ldots, 2019\}$, comprises 14,773 firm-by-event-year pairs and 5,334 unique treated firms.[23]

I use the stacked not-yet-treated DiD to avoid contamination problems that staggered treatment timing introduces in two-way fixed effects designs (Callaway and Sant'Anna 2021; Goodman-Bacon 2021; Sun and Abraham 2021). Never-treated firms represent a poor control group for treated firms. Treated and never-treated firms are on different growth paths leading up to an investment event (Appendix C.1). In contrast, treated and not-yet-treated firms are more comparable in terms of size and workforce composition immediately before each event (Table 1). Furthermore, the distribution of treated and not-yet-treated firms across sectors is similar. The pre-trends in the DiD results confirm that treated and not-yet-treated firms are on similar paths leading up to an investment event.

The baseline DiD specification compares outcomes before and after the event at treated

[23]Worker data begin in 2006, so the earliest event-year is $s = 2010$. The last observable event is in 2023, so the latest event-year is $s = 2019$.

TABLE 1. Characteristics of treated and not-yet-treated firms in pre-spike year.

Variable	(1) Treated	(2) Not-yet-treated	(3) Standardised diff.
Total employment (workers)	133.1	132.9	0.000
Early-career share (%)	13.1	11.9	0.106
Investment (thousands of euros)	139.1	93.9	0.036
Operating expenditure (thousands of euros)	36,471.7	38,113.6	−0.006
Manufacturing (%)	25.1	30.4	−0.118
Services (%)	47.5	40.9	0.133
Retail and wholesale (%)	27.4	28.7	−0.029

Notes: The table reports mean characteristics of treated and not-yet-treated firms at event time $k = -1$, the year before the investment spike. The data are from the stacked panel described in Section 3.2 with event-years s from 2010 to 2019. Units are in parentheses after each variable name and monetary variables are in thousands of 2015 euros. Column 3 reports the standardised difference in means, $(\bar{x}_T - \bar{x}_C)/\sqrt{(s_T^2 + s_C^2)/2}$, in which $\bar{x}_g$ and s_g^2 denote the mean and variance for group $g \in \{T, C\}$. Values above 0.25 indicate imbalance (Imbens and Rubin 2015). The Duncan dissimilarity index between the sectoral distributions of treated and control firms is 0.066. Number of treated firms: 5,334.

relative to not-yet-treated firms:

$$y_{jst} = \phi_{js} + \lambda_t + \theta\,(I_T \times I_P) + \epsilon_{jst}, \tag{6}$$

in which y_{jst} is the outcome for firm j in event-year s at time t. The indicator I_T equals one if j is treated in event-year s and zero if the firm is not-yet-treated. I_P represents a dummy for post-event $k \geq 0$. I add firm-by-event-year effects ϕ_{js} because firms may appear in multiple event-years. The term λ_t denotes calendar-year fixed effects. Standard errors are clustered by firm. Furthermore, I weight each observation by the corresponding firm's early-career employment at $k = -1$.[24] I also estimate a dynamic specification, in which I replace I_P with event-time dummies I_k:

$$y_{jst} = \phi_{js} + \lambda_t + \sum_{k \neq -1} \delta_k\,(I_k \times I_T) + \epsilon_{jst}. \tag{7}$$

When y_{jst} measures the EC share, the coefficient θ measures the average pre-post change in the EC share for treated compared to control firms. An estimate $\hat{\theta} < 0$ implies that firms substitute away from early-career workers after the investment spike. The estimates $\hat{\delta}_k$ trace out the dynamic path of the share around the spike.

To test for differential effects across groups, I construct another dataset stacked at the

[24] Unweighted, the estimates describe the average firm; weighted, the average early-career employer. Appendix C.6 reports unweighted estimates and estimates weighted by total employment.

firm-event-year-group-e level, rather than the firm-event-year level as in the baseline. Each observation is the share of group-e early-career workers in the firm's total employment. Then, I estimate a triple-difference specification:

$$y_{ejst} = \phi_{ejs} + \lambda_{et} + \beta \ (I_T \times I_P \times I_e) + u_{ejst}, \tag{8}$$

in which $I_e = 1$ for the college-educated, DT-exposed, or within-two-years-of-schooling group within firm j at time t in event-year s. The fixed effects ϕ_{ejs} absorb the level difference between groups, so β isolates the differential post-event decline within the same firm. All other terms are as in Equation (6).

3.3. Threats to Identification

The estimand in the stacked design is the average effect of an investment event on the EC share in treated firms relative to not-yet-treated firms. The estimate $\hat{\theta}$ identifies the effect under two conditions. The first is parallel trends: absent the investment event, the EC share would have evolved similarly in treated and control firms. The assumption cannot be tested directly, but I provide evidence for it with standard falsification tests that examine the pre-event coefficients (Figure 4).

The second assumption is no-anticipation: firms do not adjust their EC share ahead of the event. Investment spikes are an endogenous firm choice; however, identification in the stacked design rests on the *timing* of the event, not on the event itself. I find no evidence of anticipatory adjustment in my setting. Anticipatory adjustment would appear at $k = -1$, which is the omitted period in Equation (7). When I re-normalise the dynamic specification to $k = -3$ or $k = -4$ instead of $k = -1$, the estimate $\hat{\delta}_{k=-1}$ is small and not statistically different from zero (Appendix C.2). Moreover, using the same event definition and a closely related sample, Bessen et al. (2025) show the timing of firms' event-year is quasi-random and not explained by firm observables.[25]

Any remaining confounders must coincide with the timing of the investment event. As I report below, tangible capital investment and profitability spikes, defined analogously to Equation (5), are not followed by a change in the EC share (Figure 5). Other shocks, such as demand or industry-level shocks, are likely to be neutral with respect to early-career or differently-educated workers. I present evidence below that DT spikes are not neutral (Table 3). Lastly, Appendix C.3 documents that firms' use of temporary staffing agencies does not change around the investment event, so the results I see are not driven by who appears on firms' payrolls (Bijnens et al. 2026).

[25] A logit regression with firm fixed effects and time-varying wages, employment, workforce age, tenure, and gender improves on a uniform-random prediction of the event-year by only 3%. Whether a firm *ever* has an event, by contrast, is well predicted by size, sector, and workforce composition.

TABLE 2. Early-career employment effects after investment spike.

	(1) Employment share	(2) Hires share	(3) Separations share	(4) Share of hires
Treated × Post: $\hat{\theta}$	−1.64***	−1.52***	−0.11	−1.37*
	(0.57)	(0.34)	(0.35)	(0.80)
Denominator	Employment	Employment	Employment	Total hires
Units	pp	pp	pp	pp
Fixed effects				
Firm-event-year	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Observations	263k	263k	263k	230k

Notes: The table reports the pooled difference-in-differences estimate $\hat{\theta}$ from Equation (6), estimated by ordinary least squares. The outcome in Column 1 is the early-career employment share, early-career hires as a share of total employment in Column 2, early-career stayers (the complement of separations) from the previous calendar year as a share of total employment in Column 3, and the early-career share of all new hires in Column 4. All outcomes are in percentage points (0–100) and the denominator of each share is listed in the table. The sample is the stacked panel of treated and not-yet-treated firms described in Section 3.2, with event-years s from 2010 to 2019 and event time $k \in \{-4, \ldots, 3\}$. Observations are weighted by the firm's early-career employment at $k = -1$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Taken together, the above checks and arguments for the parallel trends and no-anticipation assumptions support the validity of the identification strategy used in this paper. Similar designs are used in other studies of firm investment in machinery, robots, and other automation technologies (Feigenbaum and Gross 2024; Deng et al. 2024; Aghion et al. 2025; Bisio et al. 2025; Bijnens et al. 2026).

3.4. Results

Column 1 of Table 2 reports the baseline estimate for the EC employment share. Following an investment event, the EC share decreases by 1.6 pp. The estimate represents a 12.5% decline against a pre-event mean share of 13.1%. The negative sign implies that firms substitute away from early-career workers after investing in DT.

Columns 2 and 3 of Table 2 split the numerator of the EC share into new hires and stayers (the complement of separations). Reduced inflows account for almost all of the overall effect and there is no statistically significant change for the early-career separations share.[26] Moreover, in Column 4 the early-career share of all new hires falls by 1.4 pp following the spike. Hiring

[26] In Appendix C.4, I show that the EC share falls due to a decline in early-career employment, rather than an increase in the number of other workers.

TABLE 3. Differential effects by education, field of study, and years since schooling.

	(1) College-educated	(2) Exposed field	(3) Within 2 years after schooling
Treated × Post × Group: $\hat{\beta}$	−0.95*	−0.92***	−1.22***
	(0.50)	(0.33)	(0.41)
Units	pp	pp	pp
Fixed effects			
Firm-event-year-group	Yes	Yes	Yes
Year-group	Yes	Yes	Yes
Observations	370k	376k	396k

Notes: The table reports the triple-difference estimate $\hat{\beta}$ from Equation (8), estimated by ordinary least squares. The coefficient measures the differential post-spike change in the early-career employment share of the indicated group relative to the complementary group within the same firm. Column 1 shows the differential change for college-educated relative to non-college-educated workers, Column 2 for fields of study exposed to digital technologies relative to non-exposed fields (mapping in Appendix A.4), and Column 3 for workers within two years after schooling relative to workers three to five years out. All outcomes are in percentage points (0–100). The sample is the stacked panel described in Section 3.2, with event-years s from 2010 to 2019. Columns 1 to 3 stack observations at the firm-by-group level, so the number of observations exceeds the baseline specification in Equation (6). Observations are weighted by the firm's early-career employment at $k = -1$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

therefore also tilts away from early-career workers.[27]

Table 3 displays the differential effect for various groups. The EC share decline is larger by 0.95 pp for college graduates (Column 1), by 0.92 pp for DT-exposed fields of study (Column 2), and by 1.22 pp for workers within two years after schooling (Column 3). These results are consistent with the heterogeneity in the employment cohort trends in Section 2.2 and Figure 1.[28]

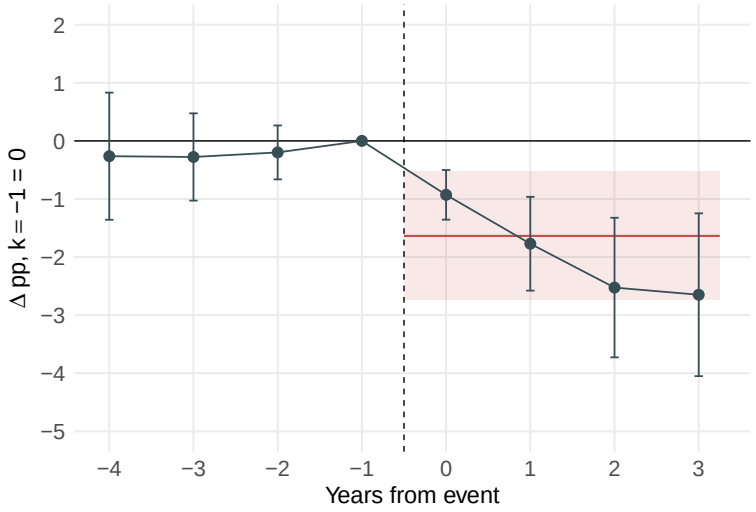
Dynamic path. Figure 4 plots the estimates $\hat{\delta}_k$ from the dynamic specification in Equation (7). I overlay the pooled DiD estimate $\hat{\theta}$ on the post-event window. Immediately after the event the estimates turn negative and significant. The decline persists and reaches −2.6 pp at $k = 3$. If I extend the event window until $k = 5$, the EC share stays depressed for the entire post-event period.

The pre-event coefficients are small and statistically insignificant at the 5% level, which

[27] I find no evidence that the EC share decline differs significantly between firms with higher cumulative profits in the pre-event period (Appendix C.4), which points away from a resource constraint argument for the decline in the EC share.

[28] The results also persist when I take the logarithm of the outcome share to control for differences in the relative employment across groups before the spike.

FIGURE 4. Early-career employment share around investment spike.



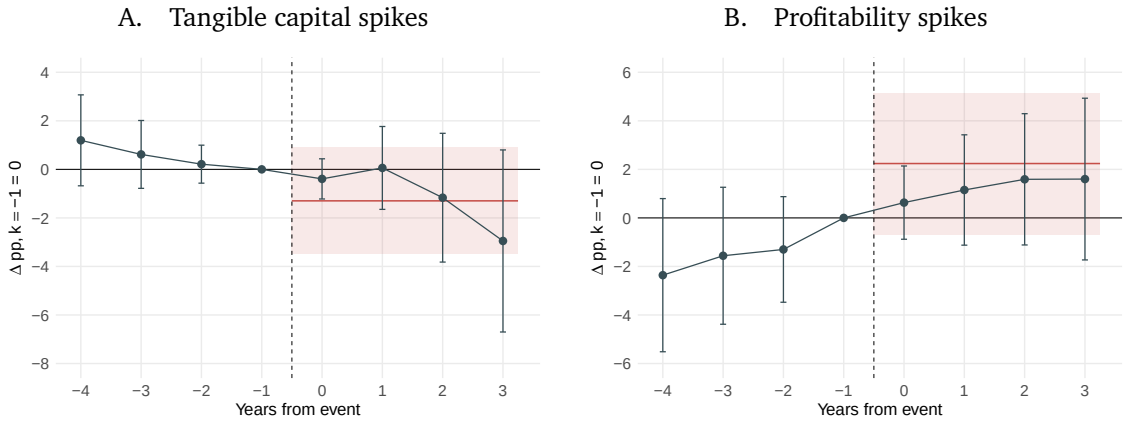
Notes: The figure plots the dynamic difference-in-differences coefficients $\hat{\delta}_k$ from Equation (7) by event time $k = t - s$. Each estimate is relative to the omitted year $k = -1$. The outcome is the early-career employment share in percentage points (0–100). The pooled estimate $\hat{\theta}$ from Equation (6) is overlaid on the post-event window with its 95% confidence interval. The sample is the stacked panel of treated and not-yet-treated firms described in Section 3.2, with event-years s from 2010 to 2019. Observations are weighted by the firm's early-career employment at $k = -1$. Bars show 95% confidence intervals with standard errors clustered at the firm level. Number of observations: 263k.

supports the parallel trends assumption. A Wald test cannot reject the null that pre-event coefficients are jointly zero ($p = 0.81$). Moreover, I find a breakdown value of $\bar{M} = 1.4$, which indicates that the pooled estimate $\hat{\theta}$ remains significant at the 5% level unless post-event violations of parallel trends exceed 1.4 times the largest pre-event violation (Rambachan and Roth 2023).

Placebo events. A natural objection to the event-study design is that DT investment spikes are correlated with other contemporaneous shocks to the firm. To address the concern, I define placebo events for tangible capital investment spikes and profitability spikes using the same definition as in Equation (5). I replace I_{js} in Equation (5) with capital investment or operating profits and re-stack the data for each event type. I plot the dynamic and pooled DiD estimates in Figure 5. Neither placebo produces an immediate and sustained decline in the EC share, which suggests that my results are specific to DT, rather than generic lumpy investment or firm growth more broadly.

Furthermore, the baseline results are robust to removing firms with DT investment spikes that ever have a spike in tangible capital investment, profitability, materials, or outsourcing during the sample period. The results are reported in Appendix C.5.

FIGURE 5. Early-career employment share around tangible investment and profit spikes.



Notes: The figure plots the dynamic difference-in-differences coefficients $\hat{\delta}_k$ from Equation (7) by event time k, relative to $k = -1$, for two placebo events. Panel A defines events as spikes in tangible-capital investment and Panel B as spikes in operating profit. Each spike is constructed as in Equation (5) with the relevant variable in place of digital-technologies investment expenditure. The outcome is the early-career employment share in percentage points (0–100). The pooled estimate $\hat{\theta}$ from Equation (6) is overlaid on the post-event window with its 95% confidence interval (tangible: −1.30 (SE : 1.12); profit: 2.24 (SE : 1.49)). Each panel uses a stacked panel of treated and not-yet-treated firms constructed as in Section 3.2, with event-years s from 2010 to 2019. Observations are weighted by the firm's early-career employment at $k = -1$. Bars show 95% confidence intervals with standard errors clustered at the firm level. Number of observations: 108k (Panel A) and 194k (Panel B).

Robustness. The baseline results hold under alternative spike definitions, fixed effects and weights, for a balanced panel and a large-firm sample, and when examining the early-career shares of hours worked (Appendix C.6). If I use the Callaway and Sant'Anna (2021) estimator in place of the stacked design to avoid contamination problems with staggered treatment timing, I find an even larger decline in the EC share following a DT investment spike. The decline appears in every sector (Appendix C.9) and is not driven by local spillovers. Controlling for the share of spiking firms in the same municipality leaves the estimate unchanged, and the EC share of not-yet-treated firms does not respond to nearby spikes (Appendix C.7). Furthermore, early-career employment at IT and business-services firms, the providers of DT services, does not rise when nearby firms spike (Appendix C.8).

4. Digital Technologies and the Returns to Early-Career Experience

The previous section shows that firms substitute away from early-career workers after investing in DT. This section asks whether the same spikes change the returns to early-career experience within firms, the second side of the entry–returns divergence. Section 4.1 describes the sample and the design. Section 4.2 presents the results.

4.1. Data and Empirical Design

Data. I combine the matched employer–employee data with firms' investment spikes. The motivating early-career fact in Section 2.3 follows each worker through potential experience $\tau \in \{0, \ldots, 5\}$, with the increase in early-career returns to experience persisting up to $\tau = 10$. Measuring the within-firm returns to early-career experience requires comparing the wages of workers at different points of their early career inside the same firm and year. For parsimony, I collapse potential experience into two windows: an *early-career* window of workers within five years of leaving school ($\tau \leq 5$) and a *post-early-career* window of workers six to ten years out ($6 \leq \tau \leq 10$). The early-career window captures the early-career individuals in the DiD design, and the post-early-career window contains workers who have just completed their early career. In this section and the quantitative model in the next section, I call the within-firm wage difference between the two windows the *experience premium*, or the *premium* for short, standing for a simplified version of the returns to early-career experience.

To observe individuals continuously until $\tau = 10$, I restrict the panel to the 2006–12 cohorts. These cohorts are also the reference group b_0 in Section 2.3, and I calibrate the quantitative model in the next section to the same cohorts. Finally, I retain worker-years at treated or not-yet-treated firms within event time $k \in \{-4, \ldots, 3\}$ for the worker's current employer. The final sample contains 875k worker-year observations.

Empirical design. Let I_H equal one if worker i is in the post-early-career window from $\tau = 6$ to $\tau = 10$ and zero otherwise. Let I_P equal one if the current employer j is past its investment spike ($k \geq 0$), and let I_k indicate event time $k = t - s$, where s denotes the event-year of firm j. I regress the log hourly wage on I_H and its interaction with I_P (pooled) or I_k (dynamic):

$$
\begin{aligned}
\text{(pooled)} \quad & \ln w_{ijt} = \psi_{jt} + \beta\, I_H + \gamma\, (I_H \times I_P) + X_{it}\delta + u_{ijt}, \\
\text{(dynamic)} \quad & \ln w_{ijt} = \psi_{jt} + \beta\, I_H + \sum_{k \neq -1} \gamma_k\, (I_H \times I_k) + X_{it}\delta + u_{ijt},
\end{aligned}
\tag{9}
$$

where ψ_{jt} are firm-year fixed effects and X_{it} contains education level-by-field, gender, and migrant-background fixed effects.[29] The coefficient β captures the within-firm experience premium, measured as the average log wage difference between post-early-career and early-career workers employed at the same firm in the same year. The coefficient γ measures the change in the premium after the spike, and γ_k traces the change in event time. Standard errors are clustered at the firm level.

Unlike the firm-level design in Section 3, Equation (9) does not stack the data by event-year. The firm-year effects ψ_{jt} restrict wage comparisons to colleagues in the same firm-year, so no cross-firm comparison is required to estimate the premium. Identification rests on the

[29] Results are robust to fully interacting these fixed effects with the firm-year effects.

TABLE 4. Within-firm experience premium around investment spike.

	(1) Firm-year	(2) Early-career	(3) Post-early-career	(4) Firm and year
Premium: $\hat{\beta}$	0.139***			
	(0.006)			
Post-event premium: $\hat{\gamma}$	0.023***	−0.015**	0.004*	0.020***
	(0.007)	(0.008)	(0.002)	(0.008)
Units	log points	log points	log points	log points
Fixed effects				
Firm-year	Yes	No	No	No
Firm-I_H	No	Yes	Yes	Yes
Year-I_H	No	Yes	Yes	Yes
Demographic	Yes	No	No	No
Demographic-I_H	No	Yes	Yes	Yes
Observations	875k	469k	430k	899k

Notes: The table reports ordinary least squares estimates of the pooled specification in Equation (9). The dependent variable is the log hourly wage. $\hat{\beta}$ is the within-firm-year experience premium comparing post-early-career workers ($5 < \tau \leq 10$) over early-career ($\tau \leq 5$) workers. The coefficient $\hat{\gamma}$ is the change in the premium after the employer's investment spike. Column 1 includes firm-by-year fixed effects. Columns 2 and 3 restrict the sample to early-career and post-early-career workers, respectively, and report the post-event wage change for each group. Because firm-by-year fixed effects would absorb the post-event indicator within a single group, these columns include firm and year fixed effects separately. Column 4 estimates the baseline with the same firm and year fixed effects, and its post-event coefficient equals the difference between Columns 3 and 2. Demographic fixed effects are education level-by-field, gender, and migrant-background. The sample is individuals from the 2006–12 cohorts (excluding the age-45 group), followed for $\tau \in \{0, \ldots, 10\}$, employed at treated or not-yet-treated firms (Section 3) within event time $k \in \{-4, \ldots, 3\}$ of their current employer. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

same two assumptions as in Section 3: absent the spike, the premium at treated firms would have followed the same path as at not-yet-treated firms, and the premium does not change in anticipation of the spike. The pre-event coefficients γ_k provide the standard falsification tests for the assumptions.

4.2. Results

Column 1 of Table 4 reports the pooled estimates $\hat{\beta}$ and $\hat{\gamma}$. The within-firm premium $\hat{\beta}$ equals 13.9 log points: post-early-career workers earn wages around 15% higher than their early-career colleagues. After the spike, the premium increases by 2.3 log points, or 18% of its pre-spike level. Together with the decline in the EC share found in Section 3, the higher premium is consistent with an entry–returns divergence inside firms following investment spikes.

A higher premium can in principle reflect lower wages for early-career workers or higher

wages for post-early-career workers. Columns 2 and 3 of Table 4 estimate the post-spike wage change separately for each group, with firm and year fixed effects in place of ψ_{jt}, which would otherwise absorb I_P. The results show that the wages of early-career workers fall by 1.5 log points, and wages of post-early-career workers rise by 0.4 log points. The difference between the estimates in Column 2 and Column 3 is equal to the change in the premium under the same fixed effects, shown in Column 4. Thus, roughly three-quarters of the increase arises from lower wages for early-career workers at $\tau \leq 5$.

Table 5 reports the differential effects for various groups. The increase in the premium is concentrated among college graduates: 5.8 log points, against 0.1 log points for workers without a degree. Moreover, the increase is also 2.4 log points larger in DT-exposed fields of study, significant at the 1% level. Both signs are consistent with the heterogeneity in the returns to early-career experience trends in Section 2.3. Column 4 splits the I_H dummy into three levels for $\tau \in \{4, 5, 6\}$ (bin 2) and $\tau \in \{7, \ldots, 10\}$ (bin 3), both relative to $\tau \in \{0, \ldots, 3\}$, showing that the premium and its change after an investment spike are concave.

Dynamic path. Figure 6 plots the dynamic estimates $\hat{\gamma}_k$ for $k \in \{-4, \ldots, 3\}$. I overlay the pooled DiD estimate $\hat{\gamma}$ on the post-event window. Immediately after the spike the estimates turn positive and significant. The higher premium persists until $k = 3$, reaching 3.5 log points.

The pre-event coefficients are small and statistically insignificant at the 5% level, which supports the parallel trends assumption. A Wald test cannot reject the null that pre-event coefficients are jointly zero ($p = 0.22$).

Robustness, selection, and match quality. As Section 3 establishes, firms reduce the inflow of early-career individuals following an investment spike. Moreover, the composition of early-career hires shifts away from college graduates and DT-exposed fields of study. The main specification always compares the experience premium within these cells. Furthermore, the premium would flatten if early-career hires are positively selected after the event, which means the estimate $\hat{\gamma}$ likely represents a lower bound. Likewise, the increase in the premium does not arise from the firm retaining better-performing post-early-career workers, since there is also an increase in the experience premium for newly hired workers ($+3.8^{***}$ log points).

Workers sort into firms based on unobservable match surplus, and this sorting may change around the spike. The standard remedies of including worker-firm match fixed effects or differencing outcomes within worker-firm spells (Altonji and Shakotko 1987; Topel 1991; Snell et al. 2018) do not suit my design. The only identifying variation with match effects included is wage growth, so $\hat{\gamma}$ would identify the difference in experience premium across bins ($I_H = 0, 1$), whereas the object of interest is the premium itself.[30] To proxy for match quality without

[30] Match effects also condition on stayers and on being hired prior to the spike to identify wage effects within worker-firm matches. I do examine how wage growth changes in Table C10, Appendix C.10.

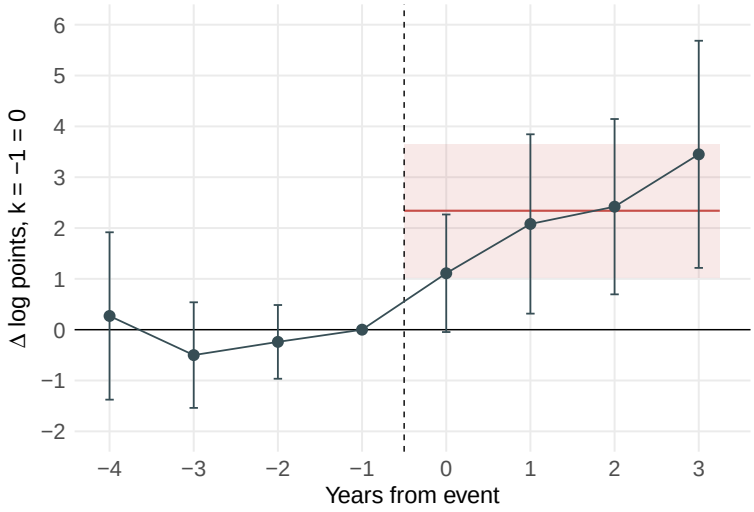
TABLE 5. Differential effects by education, field of study, and potential experience.

	(1) Baseline	(2) College-educated	(3) Exposed field	(4) Concave
Premium: $\hat{\beta}$	0.139***	0.135***	0.136***	
	(0.006)	(0.006)	(0.006)	
Post-event premium: $\hat{\gamma}$	0.023***	0.001	0.016**	
	(0.007)	(0.007)	(0.007)	
Post-event premium × group: $\hat{\gamma}^G$		0.057***	0.024***	
		(0.004)	(0.004)	
Premium, bin 2: $\hat{\beta}_2$				0.199***
				(0.011)
Premium, bin 3: $\hat{\beta}_3$				0.282***
				(0.014)
Post-event premium, bin 2: $\hat{\gamma}_2$				0.035***
				(0.012)
Post-event premium, bin 3: $\hat{\gamma}_3$				0.041***
				(0.014)
Units	log points	log points	log points	log points
Fixed effects				
Firm-year	Yes	No	No	Yes
Firm-year × group	No	Yes	Yes	No
Demographic	Yes	No	No	Yes
Demographic × group	No	Yes	Yes	No
Observations	875k	847k	844k	875k

Notes: The table reports ordinary least squares estimates of the pooled specification in Equation (9) and several extensions. The dependent variable is the log hourly wage. Column 1 repeats the baseline from Table 4. Columns 2 and 3 add the triple interaction $I_H \times I_P \times G_i$, in which G_i indicates a college degree (Column 2) or a field of study exposed to digital technologies (Column 3; Appendix A.4). Column 4 replaces the single indicator I_H with indicators for $\tau \in \{4,5,6\}$ (bin 2) and $\tau \in \{7,\ldots,10\}$ (bin 3), relative to $\tau \in \{0,\ldots,3\}$. Demographic fixed effects are education level-by-field, gender, and migrant-background. The sample is individuals from the 2006–12 cohorts (excluding the age-45 group), followed for $\tau \in \{0,\ldots,10\}$, employed at treated or not-yet-treated firms (Section 3) within event time $k \in \{-4,\ldots,3\}$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

adding match effects, I sequentially control for (i) the log hiring wage (continuous), (ii) the percentile (continuous) or vigintile (fixed effect) of the starting ($\tau = 0$) wage within i's cohort, and (iii) the within-firm rank of the hiring wage among the hiring cohort (continuous). Across all specifications, the estimates suggest that the within-firm premium increases following an investment spike (Table C10, Appendix C.10).

FIGURE 6. Path of within-firm experience premium around investment spike.



Notes: The figure plots the dynamic coefficients $\hat{\gamma}_k$ from Equation (9) by event time k of the worker's current employer, relative to the omitted year $k = -1$. The outcome is the log hourly wage, so estimates are the change in the within-firm-year experience premium in log points. The pooled estimate $\hat{\gamma}$ from the corresponding pooled regression is overlaid on the post-event window with its 95% confidence interval. The specification includes firm-by-year fixed effects and demographic fixed effects for education level-by-field, gender, and migrant-background. The sample is individuals from the 2006–12 cohorts (excluding the age-45 group), followed for $\tau \in \{0, \ldots, 10\}$, employed at treated or not-yet-treated firms (Section 3) within event time $k \in \{-4, \ldots, 3\}$. Bars show 95% confidence intervals with standard errors clustered at the firm level. Number of observations: 875k.

5. Quantitative Model

The previous sections have documented an entry–returns divergence in the aggregate, across school-leaving cohorts, and within firms, following DT investment spikes. The reduced-form DiD evidence cannot directly quantify the aggregate contribution of the rise of DT to the cohort trends or whether these changes primarily reflect substitution or augmentation channels. In this section, I build and calibrate a quantitative heterogeneous-firm model to speak to these margins. The model builds on the skill-biased technological change frameworks from Katz and Murphy (1992) and Krusell et al. (2000), adapted to DT and labour-market experience.

Environment. Time is discrete, indexed by t, and the horizon is infinite. In line with the worker DiD evidence, there are two types of labour: early-career labour, denoted by E, and post-early-career labour, denoted by H. Early-career labour refers to early-career workers within five years of schooling ($\tau \leq 5$), whereas post-early-career labour represents workers between six and ten years out ($6 \leq \tau \leq 10$), after the early career, as in Section 4. I treat all other labour as a fixed input absorbed by the span of control.[31]

[31] The denominator in the outcome variable for the firm-level DiD results is all non-early-career workers. Repeating the DiD exercise with the early-career share of early-career plus post-early-career ($6 \leq \tau \leq 10$) employment returns

A fixed mass of firms, indexed by j, produces output by combining DT with the two types of labour. There is no entry or exit of firms. All shocks in the model are idiosyncratic, so there is no aggregate uncertainty, and firms discount the future at the representative household's discount factor β.

5.1. Household

The household is populated by two types of members, early-career and post-early-career. In each period t, the household chooses consumption C_t and, for each firm j, how much labour of each type to supply: n_{Ejt} early-career labour and n_{Hjt} post-early-career labour. The household maximises expected, discounted lifetime utility:

$$\max_{\{C_t, n_{Ejt}, n_{Hjt}\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\gamma}}{1-\gamma} - \sum_{i \in \{E,H\}} \kappa_i \int \frac{n_{ijt}^{1+1/\varphi_i}}{1+1/\varphi_i} dj \right], \tag{10}$$

subject to the budget constraint:

$$C_t = \sum_{i \in \{E,H\}} \int w_{ijt} n_{ijt} dj + D_t, \tag{11}$$

in which the integral $\int dj$ runs over the mass of firms, w_{ijt} denotes the wage firm j pays to type i, and D_t stands for aggregate dividends. κ_i represents a labour supply shifter and φ_i the Frisch elasticity, and both are type-i-specific.

Labour supply. The disutility of labour supply in Equation (10) is separable across firms and types, so the household's first-order conditions give one inverse supply curve per firm-type pair:

$$w_{ijt} = \kappa_i \Lambda_t n_{ijt}^{1/\varphi_i}, \qquad \Lambda_t \equiv C_t^{\gamma}, \qquad i \in \{E,H\}. \tag{12}$$

Each firm faces an upward-sloping labour supply curve, whose position depends on the aggregate wedge Λ_t, the inverse of the marginal utility of consumption. In the baseline, I assume household utility is linear in consumption ($\gamma = 0$) and set $\kappa_i = 1$. The choice eliminates wealth effects in labour supply, so each labour market clears firm-by-firm. I make this choice to obtain firm-labour-type-specific wages that can be used to match the DiD evidence. Greenwood–Hercowitz–Huffman preferences also eliminate wealth effects from labour supply (Greenwood et al. 1988). In an extension in Appendix D.6, I allow for wealth effects and show that the baseline results are quantitatively similar. In the baseline, the labour supply curves take the

an almost identical estimate $\hat{\theta} = -1.60^{**}(0.68)$ (baseline: $\hat{\theta} = -1.64^{***}(0.57)$).

form: $w_{ijt} = n_{ijt}^{1/\varphi_i}, \;\; i \in \{E, H\}$.[32]

5.2. Firms

Production. Firm j combines early-career labour L_{Ejt}, post-early-career labour L_{Hjt}, and a stock of DT X_{jt} to produce output Y_{jt}. Production takes a two-nest form and depends on *effective* DT $M_{jt} \equiv \varepsilon_{jt} X_{jt}$, the product of the firm's actual stock and an effectiveness draw ε_{jt}:

$$Y_{jt} = \left[(L_{Ejt} + \vartheta M_{jt})^{\sigma} + (M_{jt}^{\eta} L_{Hjt})^{\sigma} \right]^{\nu/\sigma}, \tag{13}$$

with parameters $\vartheta, \eta \geq 0$, $-1 < \sigma < 1$, $0 < \nu < 1$, and $\eta\nu < 1$. DT enter the production function in two places. In the first nest, DT *substitute* for early-career labour: each unit of M supplies ϑ efficiency units of early-career labour. In the second nest, DT *augment* post-early-career labour: the term M^{η} scales the efficiency of L_H with an elasticity of $\eta \geq 0$. A defining feature of DT, especially software, is their non-rival nature. For example, software can simultaneously be used across all inputs within a firm (Haskel and Westlake 2018; Crouzet et al. 2022). In the model, non-rivalry appears in the second nest: each unit of post-early-career labour is augmented. However, the first nest features rivalry: the output that early-career labour could have produced can instead be produced by DT. The elasticity of substitution between the two nests is given by $1/(1-\sigma)$ and the parameter $\nu < 1$ governs decreasing returns to scale.

The draw ε_{jt} captures how well a firm converts its stock into output. That is, two firms with the same stock—for example, two law firms with the same predictive software—may differ in how effectively they deploy the stock, possibly due to differences in organisational capital or management practices. Effectiveness follows an AR(1) process:

$$\ln \varepsilon_{jt+1} = \rho_{\varepsilon} \ln \varepsilon_{jt} + \omega_{jt}, \qquad \omega_{jt} \sim \mathcal{N}\left(0, \sigma_{\varepsilon}^2\right). \tag{14}$$

The parameter ρ_{ε} governs how persistent shocks are, while σ_{ε}^2 captures the variance of the innovations ω_{jt}.

Timing and investment. At the start of each period, each firm observes its stock X_{jt} and draws a new ε_{jt}. Then, firms hire L_{Ejt} and L_{Hjt}. Firms take wages as given, so each firm-type wage clears where the marginal product equals the supply curve in Equation (12).[33]

Each period firms also choose investment x_{jt} in new DT. Investment is dynamic, and there is a time-to-build lag. Given a current stock of X_{jt} and investment choice x_{jt}, the next-period

[32] n_{ijt} does not separate the extensive and intensive margin of employment. I interpret the quantity as a headcount, the same margin on which the early-career employment fact (Section 2.2) and the firm-level DiD (Section 3) are measured. Both margins move in the same direction (Column 4 of Table C2).

[33] I assume firms do not internalise the labour supply curve, meaning there is no monopsony markdown.

stock X_{jt+1} is given by the following law of motion:

$$X_{jt+1} = (1-\delta)X_{jt} + x_{jt}, \tag{15}$$

in which δ is the depreciation rate. The price of the investment good q_t is the same for every firm at time t. I also assume that investment is irreversible: $x_{jt} \geq 0$.

Effectiveness shocks combined with irreversibility generate lumpy investment, because a firm cannot disinvest after an unfavourable draw. In an extension, I show that fixed adjustment costs are an alternative source of lumpiness and generate quantitatively similar results.

Firm problem. It is convenient to write the firm's maximisation problem recursively, with z' denoting next-period values of each variable z. The state variables are the existing stock X and the current effectiveness draw ε. The Bellman equation is given by:

$$V(X,\varepsilon) = \max_{L_E, L_H, x} \; Y - w_E L_E - w_H L_H - qx + \beta \mathbb{E}\big[V(X',\varepsilon') \mid \varepsilon\big], \tag{16}$$

subject to the production function in Equation (13), laws of motion in Equations (14) and (15), and the irreversibility constraint $x \geq 0$. Since firms are price-takers in their own labour markets, the wages w_E and w_H clear where the marginal product equals the labour supply.

5.3. Investment Euler

Let $\mathsf{T} \geq 0$ denote the multiplier on the irreversibility constraint $x \geq 0$. The first-order condition for investment is given by:

$$q = \beta\, \mathbb{E}\big[V_X(X',\varepsilon') \mid \varepsilon\big] + \mathsf{T}, \qquad \mathsf{T}x = 0, \tag{17}$$

in which $V_X(X',\varepsilon')$ denotes the derivative of the value function in the next period with respect to the actual stock X. From the envelope condition, V_X is:

$$V_X(X,\varepsilon) = \Pi_X(X,\varepsilon) + (1-\delta)\,\beta\, \mathbb{E}\big[V_X(X',\varepsilon') \mid \varepsilon\big] = \Pi_X(X,\varepsilon) + (1-\delta)\,(q - \mathsf{T}), \tag{18}$$

in which $\Pi(X,\varepsilon) \equiv \max_{L_E, L_H} \; Y - w_E L_E - w_H L_H$, Π_X is the derivative, and the second equality substitutes Equation (17). Rolling Equation (18) forward one period and plugging into Equation (17) yields the investment Euler equation:

$$q - \mathsf{T} = \beta\, \mathbb{E}\Big[\, \Pi_X(X',\varepsilon') + (1-\delta)\,(q' - \mathsf{T}') \,\Big|\, \varepsilon\Big]. \tag{19}$$

The left-hand side denotes the marginal cost of a unit investment today, net of the shadow value of the irreversibility constraint. The right-hand side represents the discounted marginal

benefit, equal to the marginal operating profit the additional unit generates next period, plus the value of the fraction $1 - \delta$ that survives into the following period.

When the irreversibility constraint does not bind ($T = T' = 0$) the Euler equation collapses to a user-cost condition:

$$\beta\, \mathbb{E}\big[\Pi_X(X', \varepsilon') \mid \varepsilon\big] = \big[1 - \beta(1-\delta)\big] q \;\equiv\; u(q). \tag{20}$$

Each firm invests until the user cost $u(q)$ equals the expected marginal profit of the stock. Thus, a decline in the investment price q raises firms' investment through balancing the investment Euler, since the marginal benefit is declining in the stock.

5.4. Equilibrium, Relative Labour Demand, and the Experience Premium

Combining the marginal products of L_E and L_H with the $w_i = MP_i$ condition generates a relative demand schedule for each firm, which also measures the experience wage premium in the model:

$$\ln\left(\frac{w_H}{w_E}\right) = \underbrace{\eta\sigma \ln M}_{\text{augmentation}} \underbrace{- (\sigma - 1) \ln\left(1 + \frac{\vartheta M}{L_E}\right)}_{\text{substitution}} + \underbrace{(\sigma - 1) \ln\left(\frac{L_H}{L_E}\right)}_{\text{relative scarcity}}. \tag{21}$$

The first term in the demand schedule says that a larger stock M raises the marginal product of post-early-career labour with elasticity $\eta\sigma$ (augmentation). The second term shows that an increase in the stock M lowers the marginal product of early-career labour, since each unit of M displaces ϑ units of L_E (substitution). An increase in DT ($M \uparrow$) shifts relative labour demand towards post-early-career labour through both channels. As augmentation and substitution shift labour demand away from early-career labour, the relative scarcity of early-career labour L_H/L_E increases, and the relative wage compresses at a rate $\sigma - 1$.

Appendix D.1 defines dividends, the equilibrium, and discusses the modelling choices.

5.5. Calibration

I calibrate the model to a stationary equilibrium that represents the Dutch economy over 2006–12. The period matches the reference cohorts in Section 2.3 and the worker-level DiD evidence in Section 4. Worker-level moments use the 2006–12 cohorts. Firm-level moments use the 2006–12 calendar years. I also use an aggregate DT time series from the EU-KLEMS database (Bontadini et al. 2023).

External parameters. Panel A of Table 6 lists the externally set parameters. One model period corresponds to one year. The discount factor $\beta = 0.99$ matches the average Dutch real

TABLE 6. Parameters for the baseline calibration.

Parameter	Description	Value	Source
Panel A: Externally set			
β	Discount factor	0.99	Real interest rate ≈ 1%
ν	Span of control	0.85	Winberry (2021)
δ	Depreciation rate	0.20	Peters and Taylor (2017)
φ_H	Post-early-career labour supply elasticity	2	Sokolova and Sorensen (2021)
φ_E	Early-career labour supply elasticity	2	
Panel B: Internally calibrated			
ϑ	Early-career labour substitution	1.749	Simulated method of moments
η	Post-early-career labour augmentation	0.003	Simulated method of moments
ρ_ε	Persistence of effectiveness	0.337	Simulated method of moments
σ^2_ε	Variance of effectiveness innovations	0.009	Simulated method of moments
q	Price of investment good	7.932	Simulated method of moments
σ	CES exponent between nests	0.694	Simulated method of moments

Notes: This table shows the parameters for the baseline calibration. Panel A lists externally set parameters and Panel B lists parameters estimated jointly by simulated method of moments (SMM), as described in the text. One model period coincides with a year.

interest rate over 2006–12 (World Bank). The span of control is $\nu = 0.85$ (Winberry 2021). The depreciation rate $\delta = 0.20$ lies within the range of intangible capital depreciation rates in Peters and Taylor (2017), from about 10% in pharmaceuticals to 40% for computers. Estimates of firm-level labour supply elasticities generally lie between one and three (Sokolova and Sorensen 2021; Patel 2026), but none distinguish by potential experience or age.[34] I therefore consider a 3×3 grid of geometric means $\bar{\varphi} \equiv \sqrt{\varphi_E \varphi_H} \in \{1, 2, 3\}$ and ratios $r \equiv \varphi_E/\varphi_H \in \{0.5, 1, 2\}$. I choose $\bar{\varphi} = 2$ and $r = 1$, so $\varphi_E = \varphi_H = 2$, to stay agnostic about the elasticities in the baseline. The quantitative results are similar across the grid and other values of β and δ (Appendix D.5).

Internally estimated parameters. I estimate the remaining six parameters, $\Theta \equiv \{\vartheta, \eta, \rho_\varepsilon, \sigma^2_\varepsilon, q, \sigma\}$, by simulated method of moments (SMM). For a candidate $\tilde{\Theta}$, I solve Equation (16) by value function iteration on a discretised (X, ε) grid. Then, I simulate a panel of $N = 5{,}000$ firms for $T = 35$ periods and compute the model moments. I update $\tilde{\Theta}$ with the Levenberg–Marquardt algorithm and iterate until the simulated moments converge to the empirical moments. The empirical and model-generated moments are shown in Table 7. I employ a multi-start algorithm to ensure that the results are not driven by a single starting point.

The first two moments are the main post-event DiD estimates. The 'Δ early-career employment share' moment refers to the change in the EC employment share after an investment

[34] Individual-level elasticities are usually below one (Chetty et al. 2011).

TABLE 7. Empirical targets for SMM calibration.

Moment	Source	Data	Model
DiD results			
Δ early-career employment share	Col. 1 of Tab. 2, Sec. 3	−1.640	−1.410
Δ post-event returns	Col. 1 of Tab. 4, Sec. 4	0.023	0.029
Aggregate targets			
Experience premium	Eqn. (22) coefficient, Sec. 5	0.153	0.125
Zero investment (share, 0–1)	2006–12 firm panel	0.295	0.260
DT stock / GDP	Bontadini et al. (2023)	0.148	0.149
Elasticity of subs. across age groups	Card and Lemieux (2001)	4.926	4.603

Notes: Panel 'DiD results' shows the empirical moments from the DiD evidence in Sections 3 and 4. Panel 'Aggregate targets' shows other moments. The 'Experience premium' moment denotes the within-firm experience premium estimated in Equation (22). 'Zero investment' refers to the share of all firm-year observations between 2006 and 2012 with zero recorded DT investment. The 'DT stock / GDP' ratio moment is calculated using the 'broad' measure from Bontadini et al. (2023), which includes software, databases, and organisational capital. The 'Elasticity of subs across age groups' moment refers to the elasticity of substitution ('subs.') of workers between age groups estimated in Card and Lemieux (2001). Model moments are computed using a simulated panel of firms, with investment events defined exactly as in Equation (5). The SMM estimation procedure is described in more detail in the text. The model is just-identified.

spike (Table 2, Column 1).[35] The 'Δ post-event returns' moment refers to the change in the post-early-career wage return following an investment spike (Table 4, Column 1). I compute both on the simulated panel with the stacked DiD design.[36]

The third moment is the within-firm experience premium π, estimated on the 2006–12 cohorts. I estimate the aggregate within-firm experience premium from a regression of log wages on the post-early-career worker dummy I_H, equal to one if $6 \leq \tau \leq 10$ and zero for $\tau \leq 5$

$$\ln w_{ijt} = \psi_{jt} + \pi I_H + X_{it}\delta + u_{ijt}, \tag{22}$$

where ψ_{jt} are firm-year fixed effects and X_{it} contains education level-by-field, gender, and migrant background fixed effects. The regression in Equation (22) takes the same form as the worker DiD regression in Equation (9) without the post-event interaction. I estimate $\hat{\pi} = 0.153$ (s.e. 0.0004), which implies a within-firm experience premium of 16.5%. I weight the model-implied moment by employment because Equation (22) is estimated at the worker level.

[35] The denominator in the EC share results is all non-early-career workers, who are not included explicitly in the model. Repeating the DiD exercise with the early-career share of early-career plus post-early-career ($6 \leq \tau \leq 10$) employment returns an almost identical estimate $\hat{\theta} = -1.60^{**}(0.68)$ (baseline: $\hat{\theta} = -1.64^{***}(0.57)$).

[36] Within the SMM loop, I use the Callaway and Sant'Anna (2021) estimator to reduce computation time. The two estimators agree closely, and Table 7 reports the final stacked estimates. The worker-level return regression in Table 4 includes firm-year fixed effects, so identification arises from comparing post-early-career to early-career wages within the same firm-year, which the simulated firm panel delivers directly.

Two moments discipline the firms' investment technology. First, I target a zero investment share of 29.5% from the 2006–12 firm panel. Second, I target the DT stock to GDP ratio from Bontadini et al. (2023).

The final moment is the elasticity of substitution across age groups of 4.9 from Card and Lemieux (2001). In the model, this elasticity corresponds to the *direct* elasticity of substitution across the labour types. The moment is pinned down by the production function exponent σ, which governs the elasticity between the two nests.[37]

Model fit and parameter estimates. Table 7 reports the fit. The model matches the decline in the EC share (−1.41 versus −1.64 in the data) and the increase in the returns to early-career experience (0.029 versus 0.023) following an investment spike. Moreover, the model matches the return (0.125 versus 0.153), the zero-investment share (0.260 versus 0.295), and the stock-to-GDP ratio (0.149 versus 0.148). The implied elasticity of substitution is 4.60, against 4.93 in Card and Lemieux (2001). Appendix D.2 reports untargeted moments.

The parameter estimates are shown in Panel B of Table 6. The substitution parameter $\hat{\vartheta} = 1.749$ implies that one unit of effective DT stock substitutes for 1.75 units of early-career labour. I estimate an augmentation parameter of $\hat{\eta} = 0.003$, which implies that the marginal product of post-early-career labour increases by $\hat{\eta} \times \hat{\sigma} = 0.02\%$ following a 10% increase in the effective stock (Equation (21)).[38] Effectiveness shocks are less persistent and less volatile ($\hat{\rho}_\varepsilon = 0.337$, $\hat{\sigma}^2_\varepsilon = 0.009$) than standard Hicks-neutral productivity processes, for example 0.90 and 0.23 in Winberry (2021). The investment price $\hat{q} = 7.932$ pins down the investment intensity at the steady state, since firms invest until the marginal benefit equals the marginal cost q, from Equation (20). Finally, $\hat{\sigma} = 0.694$ implies an elasticity of substitution between the nests of $1/(1 - \hat{\sigma}) \approx 3.3$.

5.6. Counterfactual: A Decline in the Investment Price

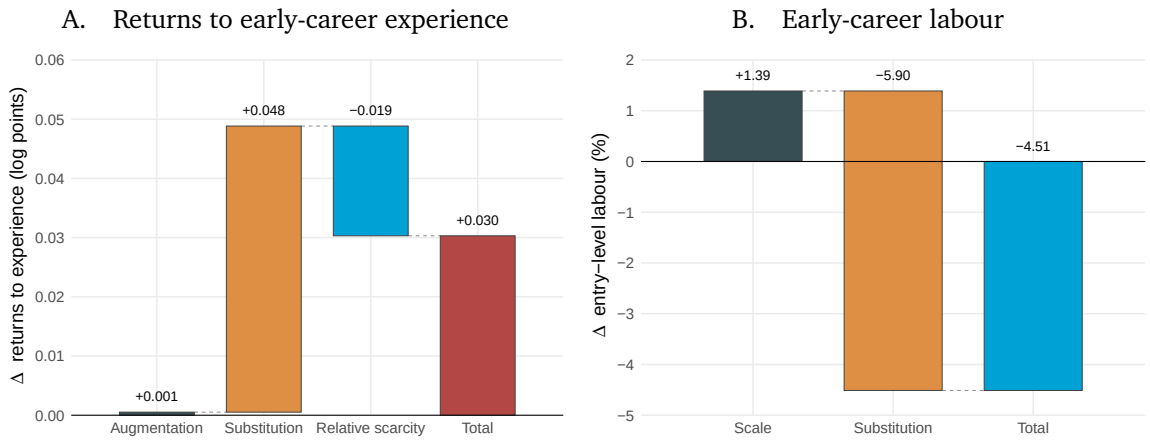
To quantify the total contribution of the rise of DT to the entry–returns divergence, I implement the following counterfactual exercise. From the EU-KLEMS data in Bontadini et al. (2023), the relative price of DT declined by 3.6% from 2006–12 to 2013–19. To arrive at this relative price, I divide the DT price deflator by the value-added price deflator. Treating the price decline as exogenous, for example due to advances in information technology, I re-solve the model at a new steady state with $q' = (1 - 0.036)\hat{q}$, holding all other parameters fixed.[39] The two steady states are meant to capture the first and last cohorts in my sample.

[37] I target the employment-weighted elasticity of substitution between the labour types. The derivation is shown in Appendix D.3.

[38] The arithmetic is: $\hat{\eta} \times \hat{\sigma} \times 0.10 = 0.003 \times 0.694 \times 0.10 \approx 0.02\%$.

[39] I do not find any evidence that the firm-level DiD estimates change over time (Appendix C.4).

FIGURE 7. Steady-state changes in early-career labour and returns to early-career experience.



Notes: This figure decomposes the counterfactual changes in the returns to early-career experience (post-early-career wage return) and early-career labour between the two steady states, following the fall in the investment price to $q' = (1 - 0.036)q$. Panel A decomposes the change in the return to experience, $\Delta \ln(w_H/w_E)$, into the augmentation, substitution, and relative scarcity channels of Equation (21). All terms are employment-weighted averages across the simulated firm distribution. Panel B decomposes the % change in early-career labour into a scale and a substitution channel. The decomposition is shown in Appendix D.4.

The price decline increases firms' investment through the user-cost Equation (20) and affects relative labour demand and the experience premium through Equation (21). By implementing the counterfactual with an exogenous price decline, the model predicts how much of the entry–returns divergence is driven just by the rise of DT.

Panel A of Figure 7 shows that the experience premium increases by 3.0 log points across steady states ('Total'). Decomposing the total effect with Equation (21), the model implies that substitution drives the increase, contributing 160% of the total, while augmentation contributes comparatively little (3%).[40] The remainder is offset because early-career labour becomes scarcer as labour demand shifts away from early-career labour due to the augmentation and substitution channels. As early-career labour becomes relatively scarcer, the relative wage becomes compressed, which offsets −63% of the increase in the experience premium.

Panel B of Figure 7 shows that early-career labour falls by 4.5% across the steady states. The decline splits into a scale and a substitution effect (see Appendix D.4). The lower investment price q' expands the scale of the economy, which increases total labour demand and raises early-career labour by 1.4%. Nevertheless, substitution lowers early-career labour by 5.9%, outweighing the positive scale effect.

Through the lens of the model, the rise of DT is quantitatively important for the entry–returns divergence that I document in the data. Acknowledging the mapping from the data to the model is imperfect, I simply divide the model-implied steady-state changes by the aggregate

[40] For reference, the median share of the experience-premium increase that augmentation accounts for around 10% across all model extensions.

cohort estimates from the data. The model accounts for 158% of the cross-cohort decline in employment from the data[41] Forces outside the model, such as the growth of flexible work arrangements and the shift towards services, may cushion the decline in the data. The model also reproduces 33% of the increase in the wage returns to early-career experience.[42] On both sides of the entry–returns divergence, substitution rather than augmentation is the dominant channel.

5.7. Model Extensions

I deliberately simplify the baseline model to isolate the substitution–augmentation asymmetry and to quantify the early-career consequences of DT with minimal machinery. Here I discuss several model extensions to show the robustness of the baseline counterfactual results. I impose the same 3.6% investment price decline for each counterfactual. Appendix D.6 reports the extensions in more detail and various other robustness checks.

Stock-to-GDP counterfactual. The baseline counterfactual compares steady-state outcomes following a 3.6% investment price decline. The price decline generates an increase in the DT stock-to-output ratio of 21.6%, from 0.15 to 0.18, whereas the EU-KLEMS data show a 22.5% increase. A price decline about -0.04 pp larger than the baseline counterfactual is needed to generate the slightly larger rise in the stock-to-output ratio. This small price difference shows that the model matches the elasticity of investment with respect to the price of investment well, even though the moment is untargeted. In the new counterfactual, the experience premium increases by 3.2 log points (baseline: 3.0) and early-career labour falls by 4.7% (baseline: 4.5%).

Wealth effects in labour supply. In the baseline I set $\gamma = 0$ in the household's utility function to eliminate wealth effects in labour supply. With $\gamma = 2$, the returns to experience rise by 3.1 log points (baseline: 3.0) and the decline in early-career labour is 5.3% (baseline: 4.5%). The decline is larger because the scale of the economy, and therefore the household's income, grows following the price decline, which reduces overall labour supply.

Fixed adjustment costs. The effectiveness shocks combined with the irreversibility of investment generate investment spikes in the baseline model. Replacing the effectiveness shocks with random fixed adjustment costs F_{jt} also generates lumpiness. In this extension, the returns

[41] From the base employment rate of 96.3% for the 2006–07 cohorts, the employment probability falls by 2.9% across cohorts (2.76/96.3; Figure 1, Section 2.2), compared with 4.51% in the model.

[42] The five-year return increases by 9.3 log points: 3.03/9.30 ≈ 0.33 (five-year return, Figure 2, Section 2.3). Taking the seven- or ten-year return would result in a higher percentage.

rise by 3.3 log points (baseline: 3.0) and the decline in early-career labour is 4.8% (baseline: 4.5%).

Experience accumulation. Labour types are fixed in the baseline model. In this extension, I allow for the stock of post-early-career workers to accumulate depending on the previous period's *employed* early-career workers. Suppose that the stock of post-early-career workers in period $t+1$ is a function of the current stock of post-early-career workers N_{Ht} and the number of employed early-career workers N_{Et}:

$$N_{Ht+1} = (1 - \delta_H)N_{Ht} + \lambda N_{Et}. \tag{23}$$

I set $\delta_H = 1/40$ for a 40-year working life and $\lambda = 1/5$ for the five-year potential experience threshold. I implement a gradual (equal per year) price decline from 2006 to 2019 as the counterfactual. By 2019, early-career labour declines by 3.49% (baseline across steady states: 4.51%) and the experience premium rises by 2.96 log points (baseline: 3.03).

6. Conclusion

This paper documents an *entry–returns divergence* in the early-career labour market. Using administrative data from the Netherlands, I show that across school-leaving cohorts the early-career employment rate fell while the returns to early-career experience rose. The divergence was more pronounced among college-educated individuals, and the rise in the returns occurred entirely within firms rather than through worker mobility or sorting across firms.

I propose firms' adoption of digital technologies as an explanation for this divergence. Digital technologies, such as software and databases, enable firms to codify their routines and knowledge, substituting for early-career workers but augmenting experience for those who do become employed. I find that, following large investment spikes, firms reduce their early-career employment share, and the returns to early-career experience increase within the firm. Both effects are stronger for college-educated individuals, mirroring the cohort trends. A heterogeneous-firm model, calibrated to the investment-spike evidence, shows that the rise of digital technologies accounts for a substantial share of the cohort trends. On both margins, substitution rather than augmentation is the dominant channel.

The early-career consequences of the rise of digital technologies have two implications that go beyond the scope of this paper. First, because experience is accumulated on the job, declining early-career employment today decreases the size of the experienced workforce tomorrow, potentially propagating into lower long-run employment and output (Asriyan et al. 2026). Second, because the entry–returns divergence was more pronounced for college-educated individuals, digital technologies may have lowered the value of a college degree *in the early-*

career labour market and weakened incentives to invest in higher education.

Finally, although the data predate generative artificial intelligence (AI), this paper offers a two-decade benchmark for AI's early-career effects. Early evidence points towards employment losses for early-career workers in AI-exposed occupations (Brynjolfsson et al. 2025; Hosseini and Lichtinger 2026). Whether AI produces a larger or smaller entry–returns divergence than the one I document for digital technologies is an important open question.

References

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67 (2): 251–333.

Acemoglu, Daron and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, vol. 4B, edited by Orley Ashenfelter and David Card, pp. 1043–1171. Elsevier.

Acemoglu, Daron and Pascual Restrepo. 2018. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108 (6): 1488–1542.

Acemoglu, Daron and Pascual Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128 (6): 2188–2244.

Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel. 2025. "Modern Manufacturing Capital, Labor Demand, and Product Market Dynamics: Evidence from France." *American Economic Journal: Macroeconomics*. Conditionally accepted.

Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. 2015. "The Skill Complementarity of Broadband Internet." *Quarterly Journal of Economics* 130 (4): 1781–1824.

Althoff, Lukas and Hugo Reichardt. 2026. "Task-Specific Technical Change and Comparative Advantage." *NBER Working Paper* (35353).

Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer. 2016. "Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success." *Journal of Labor Economics* 34 (S1): S361–S401.

Altonji, Joseph G. and Robert A. Shakotko. 1987. "Do Wages Rise with Job Seniority?" *The Review of Economic Studies* 54 (3): 437–459.

Alvehus, Johan. 2017. "Clients and Cases: Ambiguity and the Division of Labour in Professional Service Firms." *Baltic Journal of Management* 12 (4): 408–426.

Arellano-Bover, Jaime. 2024. "Career Consequences of Firm Heterogeneity for Young Workers: First Job and Firm Size." *Journal of Labor Economics* 42 (2): 549–589.

Argan, Damiano, Robert J. Gary-Bobo, and Marion Goussé. 2025. "Is There a Devaluation of Degrees? Unobserved Heterogeneity in Returns to Education and Early Experience." *IZA Discussion Paper* 18053.

Ashworth, Jared, V. Joseph Hotz, Arnaud Maurel, and Tyler Ransom. 2021. "Changes across Cohorts in Wage Returns to Schooling and Early Work Experiences." *Journal of Labor Economics* 39 (4): 931–964.

Asriyan, Vladimir, Hugo Reichardt, and Sandro Shelegia. 2026. "Immiserizing Automation." *Mimeo*.

Autor, David H. and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103 (5): 1553–1597.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113 (4): 1169–1213.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118 (4): 1279–1333.

Bagger, Jesper, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. 2014. "Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics." *American Economic Review* 104 (6): 1551–1596.

Baksy, Aniket, Daniele Caratelli, and Niklas Engbom. 2026. "The Long-term Decline of the U.S. Job Ladder." *Mimeo*.

Bartelsman, Eric, Sabien Dobbelaere, and Alessandro Zona Mattioli. 2026. "Intangible Assets and Imperfections in Product and Labor Markets." *European Economic Review* 186: 105309.

Barth, Erling, James C. Davis, Richard B. Freeman, and Kristina McElheran. 2023. "Twisting the Demand Curve: Digitalization and the Older Workforce." *Journal of Econometrics* 233 (2): 443–467.

Beaudry, Paul, David A. Green, and Benjamin M. Sand. 2016. "The Great Reversal in the Demand for Skill and Cognitive Tasks." *Journal of Labor Economics* 34 (1): 199–247.

Bessen, James, Maarten Goos, Anna Salomons, and Wiljan van den Berge. 2025. "What Happens to Workers at Firms That Automate?" *The Review of Economics and Statistics* 107 (1): 125–141.

Bianchi, Nicola and Matteo Paradisi. 2026. "Countries for Old Men: An Analysis of the Age Pay Gap." *SSRN Working Paper* (3880501).

Bijnens, Gert, Anders Humlum, and Stijn Vanormelingen. 2026. "Integrators: The Firm Boundaries of Capital-Skill Complementarity." *Mimeo*. August 2026.

Bisio, Laura, Angelo Cuzzola, Marco Grazzi, and Daniele Moschella. 2025. "The Dynamics of Automation Adoption: Firm-Level Heterogeneity and Aggregate Employment Effects." *European Economic Review* 173: 104943.

Bolhaar, Jonneke, Marloes de Graaf-Zijl, and Bas Scheer. 2018. "Three Perspectives on the Dutch Growth of Flexible Employment." *De Economist* 166 (4): 403–432.

Bontadini, Filippo, Carol Corrado, Jonathan Haskel, and Massimiliano Iommi. 2023. "Euklems and IntanProd: Industry Productivity Accounts with Intangibles." *MPRA Paper* 126218.

Borjas, George J. 1999. "The Economic Analysis of Immigration." *Handbook of Labor Economics* 3, Part A: 1697–1760. Edited by Orley C. Ashenfelter and David Card, chapter 28. Elsevier.

Brandon-Jones, Alistair, Michael Lewis, Rohit Verma, and Matthew C. Walsman. 2016. "Examining the Characteristics and Managerial Challenges of Professional Services: An Empirical Study of Management Consultancy in the Travel, Tourism and Hospitality Sector." *Journal of Operations Management* 42–43 (1): 9–24.

Bratsberg, Bernt, Erling Barth, and Oddbjørn Raaum. 2006. "Local Unemployment and the Relative Wages of Immigrants: Evidence from the Current Population Surveys." *The Review of Economics and Statistics* 88 (2): 243–263.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." *Stanford Digital Economy Lab Working Paper*.

Caicedo, Santiago, Robert E. Lucas, Jr., and Esteban Rossi-Hansberg. 2019. "Learning, Career Paths, and the Distribution of Wages." *American Economic Journal: Macroeconomics* 11 (1): 49–88.

Callaway, Brantly and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200–230.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90 (3): 414–427.

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *The Quarterly Journal of Economics* 131 (2): 633–686.

Card, David and Thomas Lemieux. 2001. "Can Falling Supply Explain the Rising Return to College for

Younger Men? A Cohort-Based Analysis." *The Quarterly Journal of Economics* 116 (2): 705–746.

Chetty, Raj, David Grusky, Maximilian Hell, Nathaniel Hendren, Robert Manduca, and Jimmy Narang. 2017. "The Fading American Dream: Trends in Absolute Income Mobility since 1940." *Science* 356 (6336): 398–406.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review* 101 (3): 471–475.

Cooper, Russell W. and John C. Haltiwanger. 2006. "On the Nature of Capital Adjustment Costs." *The Review of Economic Studies* 73 (3): 611–633.

Corblet, Pauline. 2024. "Higher-Educated Workers' Flattening Returns to Experience." *Mimeo*.

Corinth, Kevin and Jeff Larrimore. 2024. "Has Intergenerational Progress Stalled? Income Growth Over Five Generations of Americans." *Finance and Economics Discussion Series* 2024-007.

Crouzet, Nicolas, Janice C. Eberly, Andrea L. Eisfeldt, and Dimitris Papanikolaou. 2022. "The Economics of Intangible Capital." *Journal of Economic Perspectives* 36 (3): 29–52.

Daruich, Diego, Sabrina Di Addario, and Raffaele Saggio. 2023. "The Effects of Partial Employment Protection Reforms: Evidence from Italy." *The Review of Economic Studies* 90 (6): 2880–2942.

Deming, David J. 2021. "The Growing Importance of Decision-Making on the Job." *Mimeo*.

Deming, David J. 2025. "Why Do Wages Grow Faster for Educated Workers?" *Journal of Labor Economics*. Forthcoming.

Deng, Liuchun, Steffen Müller, Verena Plümpe, and Jens Stegmaier. 2024. "Robots, Occupations, and Worker Age: A Production-Unit Analysis of Employment." *European Economic Review* 170: 104881.

Elsby, Michael W. L. and Matthew D. Shapiro. 2012. "Why Does Trend Growth Affect Equilibrium Employment? A New Explanation of an Old Puzzle." *American Economic Review* 102 (4): 1378–1413.

Eurostat. 2026a. "Persons in Part-Time Employment by Sex and Age (`lfsa_eppga`)." EU Labour Force Survey, Eurostat database. https://ec.europa.eu/eurostat/databrowser/view/lfsa_eppga/default/table. Accessed August 2026.

Eurostat. 2026b. "Temporary Employees as a Percentage of the Total Number of Employees, by Sex, Age and Citizenship (`lfsa_etpgan`)." EU Labour Force Survey, Eurostat database. https://ec.europa.eu/eurostat/databrowser/view/lfsa_etpgan/default/table. Accessed August 2026.

Feigenbaum, James and Daniel P. Gross. 2024. "Answering the Call of Automation: How the Labor Market Adjusted to Mechanizing Telephone Operation." *The Quarterly Journal of Economics* 139 (3): 1879–1939.

Figueiredo, Ana, Agnieszka Markiewicz, and Elliott Weder. 2026. "Firms and the Wage Penalty of Temporary Work." *Mimeo*.

Foster, Lucia, John C. Haltiwanger, and C. J. Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." *New Developments in Productivity Analysis* pp. 303–372. Edited by Charles R. Hulten, Edwin R. Dean and Michael J. Harper. University of Chicago Press.

Goldin, Claudia and Lawrence F. Katz. 2008. *The Race between Education and Technology*. Cambridge, MA: Harvard University Press.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225 (2): 254–277.

Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased

Technological Change and Offshoring." *American Economic Review* 104 (8): 2509–2526.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman. 1988. "Investment, Capacity Utilization, and the Real Business Cycle." *The American Economic Review* 78 (3): 402–417.

Haskel, Jonathan and Stian Westlake. 2018. *Capitalism without Capital: The Rise of the Intangible Economy*. Princeton University Press.

Hopenhayn, Hugo and Richard Rogerson. 1993. "Job Turnover and Policy Evaluation: A General Equilibrium Analysis." *Journal of Political Economy* 101 (5): 915–938.

Horowitz, Joel L. and Charles F. Manski. 2000. "Nonparametric Analysis of Randomized Experiments with Missing Covariate and Outcome Data." *Journal of the American Statistical Association* 95 (449): 77–84.

Hosseini, Seyed M. and Guy Lichtinger. 2026. "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data." *Mimeo*.

Humlum, Anders and Emilie Vestergaard. 2025. "Large Language Models, Small Labor Market Effects." *NBER Working Paper* 33777. Later circulated as "Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI".

Imbens, Guido W. and Donald B. Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge: Cambridge University Press.

Jaimovich, Nir, Seth Pruitt, and Henry E. Siu. 2013. "The Demand for Youth: Explaining Age Differences in the Volatility of Hours." *American Economic Review* 103 (7): 3022–3044.

Jeong, Hyeok, Yong Kim, and Iourii Manovskii. 2015. "The Price of Experience." *American Economic Review* 105 (2): 784–815.

Kahn, Lisa B. 2010. "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy." *Labour Economics* 17 (2): 303–316.

Katz, Lawrence F. and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *The Quarterly Journal of Economics* 107 (1): 35–78.

Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante. 2000. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica* 68 (5): 1029–1053.

McGrattan, Ellen R. and Edward C. Prescott. 2010. "Unmeasured Investment and the Puzzling US Boom in the 1990s." *American Economic Journal: Macroeconomics* 2 (4): 88–123.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–1725.

Oreopoulos, Philip, Till von Wachter, and Andrew Heisz. 2012. "The Short- and Long-Term Career Effects of Graduating in a Recession." *American Economic Journal: Applied Economics* 4 (1): 1–29.

Oster, Emily. 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business & Economic Statistics* 37 (2): 187–204.

Patel, Aseem. 2026. "Labour Market Power and Automation." *Mimeo*.

Peters, Ryan H. and Lucian A. Taylor. 2017. "Intangible Capital and the Investment-*q* Relation." *Journal of Financial Economics* 123 (2): 251–272.

Rambachan, Ashesh and Jonathan Roth. 2023. "A More Credible Approach to Parallel Trends." *Review of Economic Studies* 90 (5): 2555–2591.

Rothstein, Jesse. 2023. "The Lost Generation? Labor Market Outcomes for Post-Great Recession Entrants." *Journal of Human Resources* 58 (5): 1452–1479.

Silvestrini, Riccardo and Elliott Weder. 2026. "Don't Touch My Deductions: The Macroeconomic Effects

of Balanced Corporate Tax Reforms." *Mimeo*.

Snell, Andy, Pedro Martins, Heiko Stüber, and Jonathan P. Thomas. 2018. "Bias in Returns to Tenure When Firm Wages and Employment Comove: A Quantitative Assessment and Solution." *Journal of Labor Economics* 36 (1): 47–74.

Sokolova, Anna and Todd Sorensen. 2021. "Monopsony in Labor Markets: A Meta-Analysis." *ILR Review* 74 (1): 27–55.

Statistics Netherlands. 2016. "ICT-gebruik bij bedrijven; bedrijfstak en bedrijfsgrootte, 2015." StatLine table 83248NED. https://www.cbs.nl/nl-nl/cijfers/detail/83248NED.

Sun, Liyang and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics* 225 (2): 175–199.

Topel, Robert. 1991. "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority." *Journal of Political Economy* 99 (1): 145–176.

Topel, Robert H. and Michael P. Ward. 1992. "Job Mobility and the Careers of Young Men." *The Quarterly Journal of Economics* 107 (2): 439–479.

van Bezooijen, Emiel, Wiljan van den Berge, and Anna Salomons. 2024. "The Young Bunch: Youth Minimum Wages and Labor Market Outcomes." *ILR Review* 77 (3).

von Wachter, Till. 2020. "The Persistent Effects of Initial Labor Market Conditions for Young Adults and Their Sources." *Journal of Economic Perspectives* 34 (4): 168–194.

Webb, Matthew D. 2023. "Reworking Wild Bootstrap-Based Inference for Clustered Errors." *Canadian Journal of Economics* 56 (3): 839–858.

Winberry, Thomas. 2021. "Lumpy Investment, Business Cycles, and Stimulus Policy." *American Economic Review* 111 (1): 364–396.

Appendix A. Data

This appendix reports summary statistics and documents the construction of the digital-technologies (DT) exposure measures used in the main text.

A.1. Individual-Level Summary Statistics

Table A1 reports individual-level summary statistics for the 2006 and 2019 cohorts. Recent early-career cohorts leave school later and are more likely to be female, to have a migrant background, and to hold a college degree.

Conditional on employment, hourly wages are higher for age-45 workers than for early-career workers. Hourly wage growth from $\tau = 0$ to $\tau = 1$, however, is €2.6 larger for early-career workers than for the age-45 group in the 2006 cohort. For the 2019 cohort, the gap is €3.5, suggesting that unconditional returns to early-career experience grew across the two cohorts.

TABLE A1. Summary statistics for 2006 and 2019 cohorts, early-career and age-45 groups.

	Cohort			
	2006		2019	
	Early-career	Age-45	Early-career	Age-45
Age at $\tau = 0$ (years)	21.3	45.0	21.9	45.0
Female (%)	50.5	54.4	51.0	54.2
Migrant (%)	17.2	15.9	17.7	19.2
College-educated (%)	41.5	15.5	44.5	30.0
Hourly wage at $\tau = 0$ (€)	9.4	22.2	9.5	22.0
Hourly wage at $\tau = 1$ (€)	12.4	22.6	13.3	22.3
Individuals	124k	188k	159k	143k

Notes: The table reports means for the early-career and age-45 groups in the 2006 and 2019 cohorts, as defined in Section 2.1. $\tau = t - c$ denotes potential experience, so $\tau = 0$ is the cohort year. Migrant refers to an individual with at least one foreign-born parent, following the Statistics Netherlands classification. College-educated refers to a highest qualification from an applied-science or research university. Hourly wages are conditional on employment and in 2015 euros. The final row reports the number of individuals in thousands.

A.2. Firm-Level Summary Statistics

Table A2 reports summary statistics for the PS firm panel, pooling all firm-year observations from 2000 to 2023. The average firm employs 56 workers and invests around €152k in DT per year. Around 31% of firm-years record no DT investment.

TABLE A2. Summary statistics from full firm sample.

	Percentile				Count (thousands)	
Variable	25th	50th	75th	Mean	Observations	Firms
Value added	200.7	1,609.3	6,181.7	16,278.0	862k	341k
Total workers	2	12	36	55.7	1,091k	426k
DT investment	0	10.5	45.3	151.8	601k	121k
Investment / worker	0	0.33	1.11	1.29	601k	121k
Zero investment				0.31	601k	121k

Notes: The table reports summary statistics for the Production Statistics firm panel described in Section 3.1. I pool all firm-year observations from 2000 to 2023. Value added and digital-technologies (DT) investment are in thousands of 2015 euros. Investment per worker divides digital-technologies investment by the firm's number of workers. Zero investment is the share of firm-years with no digital-technologies investment. The final two columns report the number of firm-year observations and of unique firms, in thousands, on which each row is computed.

A.3. The Rise of Digital Technologies and Early-Career Exposure

Firm-level investment. Panel A of Figure A1 plots annual DT investment per worker at the firm level, in euros relative to 2000. Total investment per worker ('Total') rose by roughly €800 from a base of €1,200, a 70% increase. The increase was driven by the upper percentiles of the firm distribution.

Early-career exposure. Conditional on employment, recent cohorts of early-career individuals have become more exposed to DT at their current employer. Let P_{ijt} equal one if individual i's employer j in year t has strictly positive DT investment in year t. I estimate the composition-adjusted probability that members of cohort group b work at an investing firm, relative to the reference cohorts $b_0 = \{2006, 2007\}$:

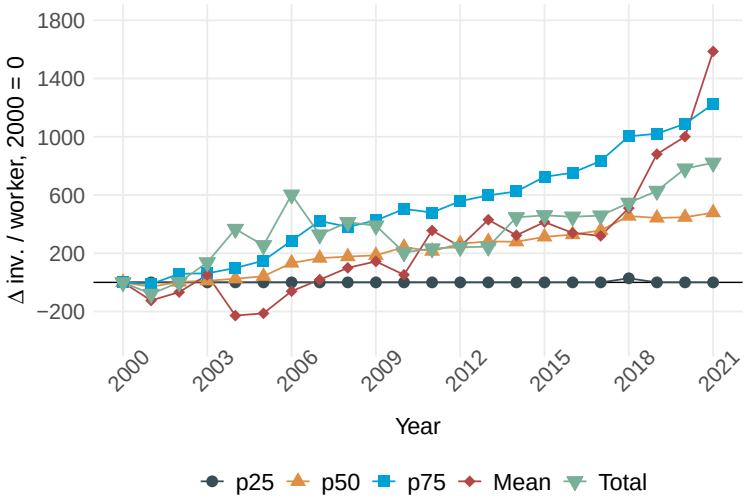
$$P_{ijt} = \sum_{b \neq b_0} \beta_b I_b + \upsilon_{ae} + \lambda_t + \zeta_g + \nu_m + u_{ijt}. \tag{A1}$$

All other terms carry the same meaning as in the employment regression in Equation (1). The sample comprises employed early-career individuals whose employer appears in the PS panel (around 50% of all employed early-career individuals). Panel B of Figure A1 plots the estimated $\hat{\beta}_b$ by cohort group, relative to the 2006–07 reference cohorts. Exposure rose by 2.6 pp over the sample, from an initial mean of 78%. The increase was comparable for college-educated and non-college-educated individuals.

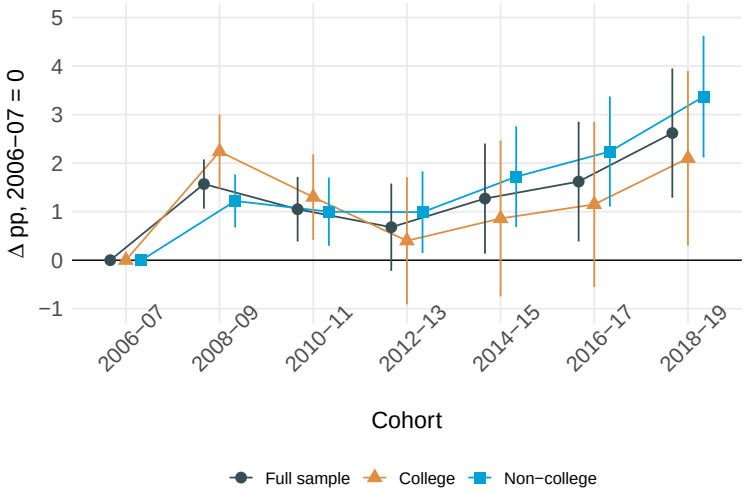
Because the administrative data do not record occupations, I turn to the Dutch labour force survey (2013–25) to understand reallocations across occupations. The share of workers aged 25 to 35 in exposed occupations rose by 9.6 pp from 2013 to 2025, from a base of 34%. The share for workers aged 35 to 55 increased by less. Appendix A.5 reports the occupation-level exposure mapping.

FIGURE A1. The rise of digital-technologies investment and early-career exposure.

A. Investment / worker, change from 2000



B. Probability of employment at an investing firm



Notes: Panel A plots the change since 2000, in 2015 euros, in annual digital-technologies investment per worker at the 25th, 50th, and 75th percentiles of the firm distribution, the unweighted mean across firms, and total investment divided by total workers ('Total'). The sample is all firm-year observations in the Production Statistics panel with non-missing investment. I include zeros. Number of observations: 345k, with 77k firms. Panel B plots the estimated cohort effects $\hat{\beta}_b$ from Equation (A1), estimated by ordinary least squares. The outcome variable is an indicator for the current employer having positive digital-technologies investment, in percentage points relative to the 2006–07 reference cohorts. I estimate Equation (A1) for the full sample and separately by education level. The sample is employed early-career individuals from Section 2.1 whose employer appears in the Production Statistics panel. Number of observations: 4.4 million. Bars show 95% confidence intervals with standard errors clustered at the cohort level.

A.4. Education Field Exposure Codes

Table A3 reports the classification of fields of study by DT exposure. I use the broad fields of the ISCED-F 2013 classification and code each field manually as exposed or not exposed. As a check on the classification, individuals in exposed fields are around 10 pp more likely to work at a DT-investing firm in a given year. The main results are robust to classifying 'Engineering, manufacturing, and construction' and 'Social sciences, journalism, and information' as Exposed, too.

TABLE A3. Classification of fields of study by digital-technologies exposure.

Field of study	Exposure
Generic programmes and qualifications	Not exposed
Education	Not exposed
Arts and humanities	Not exposed
Social sciences, journalism, and information	Not exposed
Business, administration, and law	Exposed
Natural sciences, mathematics, and statistics	Exposed
Information and communication technologies	Exposed
Engineering, manufacturing, and construction	Not exposed
Agriculture, forestry, fisheries, and veterinary	Not exposed
Health and welfare	Not exposed
Services	Not exposed
Field unknown	Not exposed

Notes: Fields of study are the broad fields of the International Standard Classification of Education: Fields of Education and Training 2013. A field is coded as exposed if its core tasks involve using, interacting with, or maintaining digital technologies. I code the mapping manually. The main results are robust to classifying 'Engineering, manufacturing, and construction' and 'Social sciences, journalism, and information' as exposed to DT, too.

A.5. Occupational Exposure Codes

Table A4 reports the classification of occupations by DT exposure. I use the 3-digit occupation segments of the 2014 Statistics Netherlands classification (BRC 2014), taken from the Statistics Netherlands occupation classification page. I manually code each occupation as exposed or not exposed.

TABLE A4. Classification of occupations by digital-technologies exposure.

Code	Occupation	Exposure
041	Business and administration professionals	Exposed
042	Business and administration associate professionals	Exposed
043	Clerical support workers	Exposed
051	Chief executives and general managers	Exposed
052	Administrative and commercial managers	Exposed
053	Production and specialised services managers	Exposed
054	Hospitality, retail and other services managers	Exposed
055	Managers not further specified	Exposed
061	Government officials and administrators	Exposed
062	Legal professionals	Exposed
071	Science and engineering professionals	Exposed
072	Science and engineering associate professionals	Exposed
081	ICT professionals	Exposed
082	ICT associate professionals	Exposed
011	Teachers	Not exposed
012	Sports instructors	Not exposed
013	Childcare workers and teaching assistants	Not exposed
021	Authors and artists	Not exposed
022	Arts and culture associate professionals	Not exposed
031	Marketing, public relations and sales advisers	Not exposed
032	Sales representatives and buyers	Not exposed
033	Sales workers	Not exposed

Continued on next page

Table A4 continued from previous page

Code	Occupation	Exposure
063	Protective services workers	Not exposed
073	Construction workers	Not exposed
074	Metal workers and machinery mechanics	Not exposed
075	Food processing and other craft workers	Not exposed
076	Electricians and electronics fitters	Not exposed
077	Machine operators and assemblers	Not exposed
078	Labourers in construction and manufacturing	Not exposed
091	Market gardeners, crop and livestock farmers	Not exposed
092	Agricultural labourers	Not exposed
101	Doctors, therapists and specialised nurses	Not exposed
102	Social science professionals	Not exposed
103	Health associate professionals	Not exposed
104	Social workers and residential care workers	Not exposed
105	Personal care workers	Not exposed
111	Personal service workers	Not exposed
112	Cleaners and kitchen helpers	Not exposed
121	Drivers and mobile plant operators	Not exposed
122	Transport and logistics labourers	Not exposed
131	Other	Not exposed

Notes: Occupations are the 41 occupation segments (beroepssegmenten) of the 2014 occupation classification of Statistics Netherlands and the Research Centre for Education and the Labour Market (Beroepenindeling ROA-CBS 2014). Occupation names are my own translations of the Dutch labels. An occupation is manually coded as exposed if its core tasks involve using, interacting with, or maintaining digital technologies.

Appendix B. Early-Career Facts

This appendix reports robustness checks for the cohort trends in employment (Section 2.2) and in the returns to early-career experience (Section 2.3).

B.1. Ramp Adjustment of the Cohort Employment Effects

Within a cohort, employment rises in the first years after leaving school (Appendix B.3). I refer to this profile as the ramp and write it as $g(e, \tau)$, where τ is the number of years since leaving school and $g(e, 0) = 0$. The term e denotes education level and field of study, which allows $g(\cdot)$ to differ by e.

Age a equals school-leaving age s plus years since school τ, $a = s + \tau$. I include age fixed effects in the main employment regression to control for life-cycle effects, such as marriage, parenthood, or home ownership, that can affect employment. Over cohorts the school-leaving age has increased (see Appendix A.1). So, individuals at some age a' in more recent cohorts are pushed down the ramp: at the same age they have spent fewer years out of school.

To see how this affects the cohort effects, consider everyone aged a' with education level-by-field e. For an individual in this group, years since school are $\tau = a' - s$. Let $\pi_c(\tau)$ denote the share of cohort c in the group that is τ years out of school. Mean employment of cohort c in the group is then given by:

$$\bar{E}_c = \beta_c + \lambda_{e,a'} + \sum_{\tau} \pi_c(\tau)\, g(e, \tau),$$

where β_c is the cohort effect and $\lambda_{e,a'}$ the education level-by-field-by-age effect.[43] The last term is the average position of cohort c on the ramp.

The education-by-age effect is assumed to be the same for every cohort. Comparing cohort c with the reference cohort c_{ref} therefore gives

$$\bar{E}_c - \bar{E}_{\text{ref}} = (\beta_c - \beta_{\text{ref}}) + \sum_{\tau} \big[\pi_c(\tau) - \pi_{\text{ref}}(\tau)\big]\, g(e, \tau).$$

A rising school-leaving age thus shifts the shares $\pi_c(\tau)$ toward low τ, where the ramp is lowest. For $\tau > 0$, where $g(e, \tau) > 0$, the shares π_{ref} rise relative to the shares π_c, which makes the bias negative. To link this back to the cohort employment regression, note that Equation (1) contains no variable for years since school, and its age and year effects are shared across cohorts. The bias is therefore absorbed by the cohort effect.

Ramp adjustment. I remove the bias in three steps, estimating a non-linear $g(e, \tau)$.

[43] Here, I omit year and other fixed effects for parsimony.

First, I estimate the ramp, $g(e,\tau)$. Following each cohort over time traces out the ramp $g(\cdot)$, but also absorbs business cycle fluctuations. To separate the two, I estimate on early-career workers and the age-45 group together:

$$E_{it} = \alpha_{et} + (1 - I_y)\big[\mu_c + g(e,\tau_{it})\big] + u_{it}. \tag{B1}$$

Here I_y equals one for the age-45 group, α_{et} are education-by-year effects assumed to be the same for both groups, and μ_c are cohort effects for early-career workers. The age-45 group has no ramp because they are not entering the labour market, so their employment identifies α_{et} and allows the business cycle to be netted out. I include cohort fixed effects so that $g(e,\tau)$ does not absorb the cohort effects.

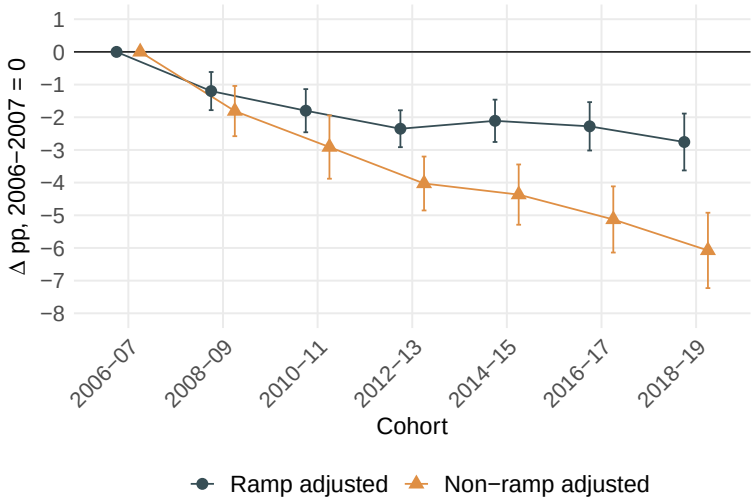
Second, I extract $\hat{g}(e,\tau)$. Third, I estimate Equation (1) with $E_{it} - \hat{g}(e,\tau)$ as the outcome. Because least squares is linear in the outcome, the cohort coefficients equal

$$\hat{\beta}_c^{\text{corr}} = \hat{\beta}_c - \hat{\beta}_c^{\text{mech}}.$$

The outcome variable equals employment with the typical (pooled) employment gain from τ netted out. A regression with $\hat{g}(e,\tau)$ as the outcome would estimate *just* the mechanical component $\hat{\beta}_c^{\text{mech}}$ of the cohort effects—if employment depended on years since school alone, and not the cohort itself.

Figure B1 compares the full-sample cohort employment effects with and without the adjustment. Without the adjustment, the 2018–19 cohorts are 6.1 pp less likely to be employed than the 2006–07 cohorts; with the adjustment, they are 2.8 pp less likely. This implies that the mechanical component of the cohort effects is 3.3 pp.

FIGURE B1. Cohort effects on early-career employment probabilities, with and without the ramp adjustment.



Notes: The figure plots the estimated cohort employment effects $\hat{\beta}_b$ from Equation (1) for the full early-career sample, estimated by ordinary least squares. 'Ramp adjusted' uses the employment indicator net of the qualification-specific employment ramp over the years since school, as in the rest of the paper. 'Non-ramp adjusted' uses the unadjusted employment indicator. Estimates are in percentage points relative to the 2006–07 cohorts. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 10.6 million.

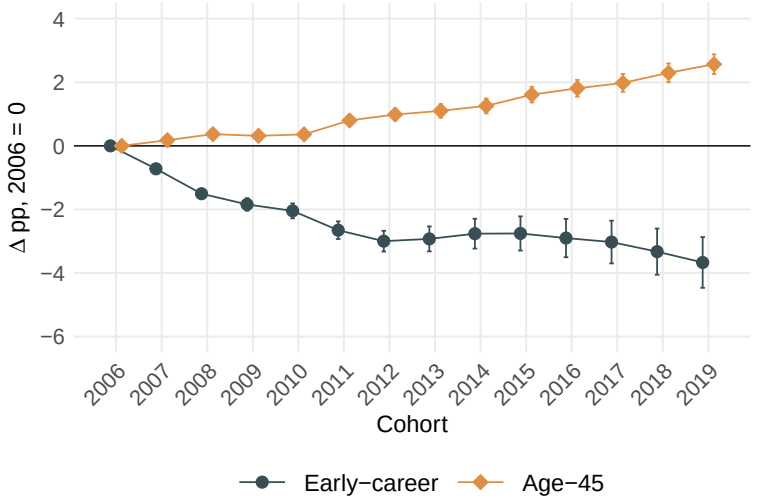
B.2. Cohort Employment Effects for Every Cohort

The baseline specification bins adjacent cohorts. Here I estimate a separate effect for every cohort c:

$$E_{it} = \sum_{c \neq c_0} \beta_c I_c + \upsilon_{ae} + \lambda_t + \zeta_g + \nu_m + u_{it}, \tag{B2}$$

in which c now replaces the adjacent cohort bin b from the baseline specification in Equation (1). The reference group c_0 is the 2006 cohort. I plot the estimates of β_c for the early-career and age-45 groups separately in Figure B2. The results are qualitatively similar to the binned specification.

FIGURE B2. Cohort effects on employment probabilities, early-career and age-45 individuals, single cohorts.



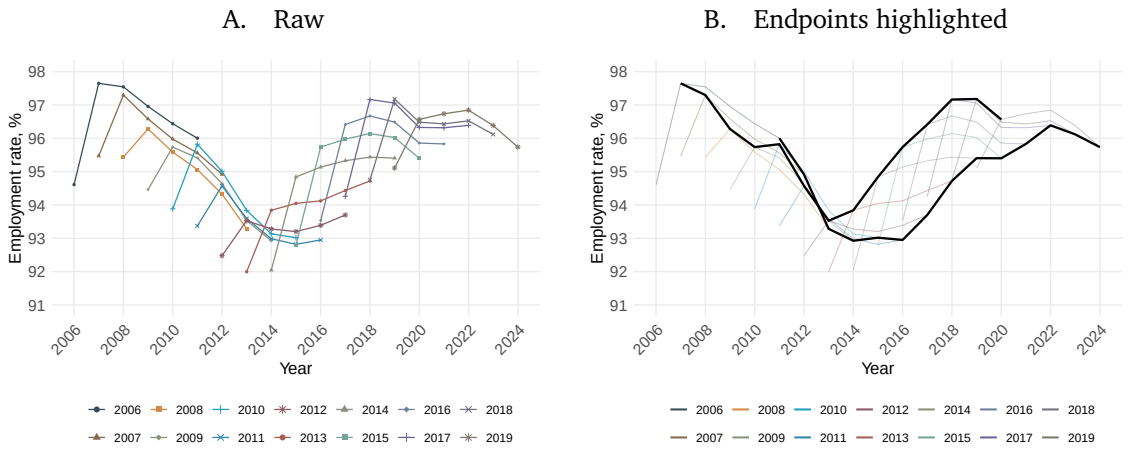
Notes: The figure plots the estimated cohort employment effects $\hat{\beta}_c$ from Equation (B2) for each single cohort c, estimated by ordinary least squares. The outcome is an indicator for being employed at any point in the year. Estimates are in percentage points relative to the 2006 cohort. The sample is the early-career and age-45 groups described in Section 2.1: the 2006–19 cohorts observed for $\tau \in \{0, \ldots, 5\}$. 'Early-career' estimates Equation (B2) on all early-career individuals, with the ramp-adjusted employment indicator (Appendix B.1); 'Age-45' estimates Equation (B2) with education level-by-field fixed effects on the age-45 group. I do not adjust the age-45 estimates for the τ-ramp, as in Appendix B.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 10.6m ('Early-career'), 9.5m ('Age-45').

B.3. Raw Employment Profiles

Panel A of Figure B3 plots raw employment rates by cohort and years since schooling. Panel B highlights the $\tau = 1$ and $\tau = 5$ values of each cohort. Both panels show a strong business cycle pattern common to all cohorts. Comparing early and late cohorts, there was a slight decline in raw employment rates. The regression in Equation (1) nets out the macroeconomic trends and other compositional changes and uncovers the sharper decline across cohorts.

From the raw profiles, we can also see the employment 'ramp' from $\tau = 0$ to $\tau = 1$ by cohort.

FIGURE B3. Raw employment rates by cohort and years since schooling.



Notes: The figure plots the mean employment rate for each cohort c by calendar year t within $\tau = t - c \in \{0, 1, \ldots, 5\}$. The sample is the early-career panel described in Section 2.1. Panel A plots every cohort-year combination; Panel B highlights the $\tau = 1$ and $\tau = 5$ values of each cohort. Number of observations: 10.6 million.

B.4. Cohort Employment Effects for Age-45 Individuals

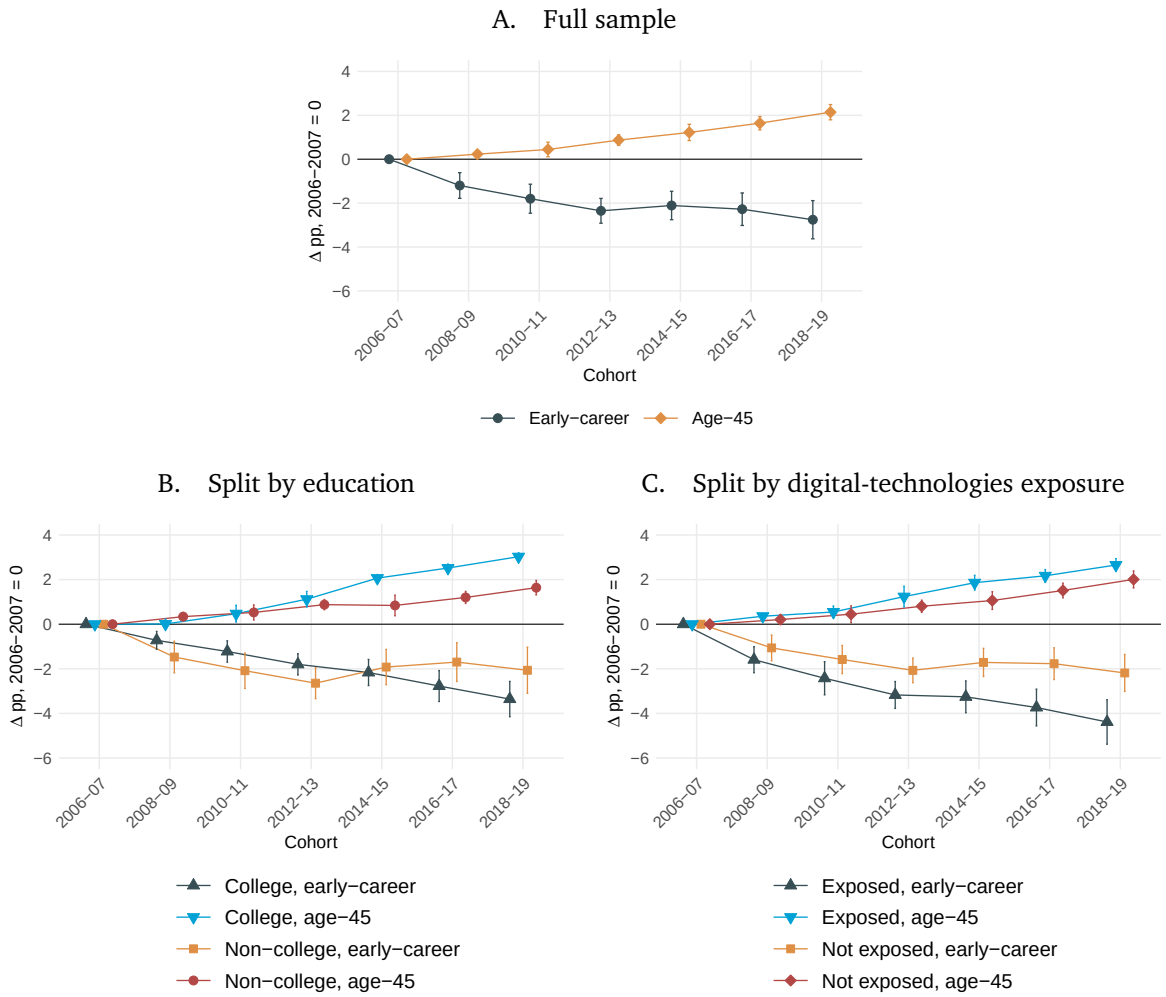
Here I estimate the cohort employment effects for the age-45 group alone. I re-estimate the baseline specification of Section 2.2:

$$E_{it} = \sum_{b \neq b_0} \beta_b I_b + \upsilon_e + \lambda_t + \zeta_g + \nu_m + u_{it}, \tag{B3}$$

on age-45 individuals only. Since all individuals are the same age by construction, I replace the education level-by-field-by-age effects with education level-by-field effects υ_e. Figure B4 plots $\hat{\beta}_b$ alongside the ramp-adjusted early-career estimates. I do not adjust the age-45 estimates for the τ-ramp, unlike the school-leaving cohorts in Appendix B.1. Panel A shows that, in contrast to the early-career decline, the employment probability of age-45 individuals rose by 2.1 pp from the 2006–07 to the 2018–19 cohorts. For reference, the unconditional employment rate of the age-45 group is around 85% for the reference cohorts.

Panel B shows that the increase was more pronounced for college graduates than for non-college graduates. Panel C demonstrates that the increase was also more pronounced for exposed than non-exposed fields of study. Both gradients have the opposite sign to the early-career gradients in Figure 1.

FIGURE B4. Cohort effects on employment probabilities, early-career and age-45 groups.



Notes: The figure plots estimated cohort employment effects $\hat{\beta}_b$ for the early-career group from Equation (1) and for the age-45 group from Equation (B3). For the age-45 estimation in Equation (B3) I drop the age fixed effects because all individuals are the same age by construction. Adjacent cohorts are binned into groups b, so each estimate is in percentage points relative to the 2006–07 reference cohorts. Panel A reports the full-sample estimates for each group, Panel B splits by education level, and Panel C by field-of-study exposure to digital technologies (mapping in Appendix A.4). Both samples are described in Section 2.1. Note, education coverage is not complete for age-45 individuals. I do not adjust the age-45 estimates for the τ-ramp, as in Appendix B.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations, early-career and age-45 respectively: full sample 10.6 and 9.5 million; college 4.5 and 3.2 million; non-college 6.1 and 6.3 million; exposed 2.7 and 2.2 million; non-exposed 7.9 and 7.3 million.

B.5. Demand- and Supply-Constrained Non-Employment

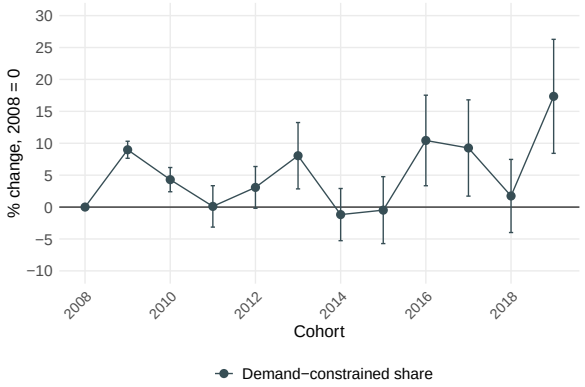
The decline in employment across cohorts documented in Section 2.2 may reflect weaker labour demand or weaker labour supply. To separate the two, I use the LISS panel, which asks non-employed respondents why they are currently not in employment. I code three answers as demand-constrained: job seeker following loss of job; school-leaver seeking a first job; and (re-)entrant seeking work. All remaining answers, such as household constraints, are coded as supply-constrained.

Then, I take all non-employed individuals in the LISS survey and let D_{it} be a dummy equal to one if individual i is in demand-constrained non-employment in year t, and zero if they are supply-constrained. I estimate the cohort profile of demand-constrained non-employment with the same specification as the employment regression in Equation (1), replacing the outcome with D_{it}:

$$D_{it} = \sum_{b \in \mathcal{B} \setminus b_0} \beta_b I_b + \upsilon_{ae} + \lambda_t + \zeta_g + \nu_m + u_{it}. \tag{B4}$$

Figure B5 plots the estimated cohort effects $\hat{\beta}_b$. The 2008 cohort is the reference cohort, since it is the earliest year in which I can define a cohort of school-leavers with the LISS panel. The estimates show that, if anything, demand-constrained non-employment became more prevalent among recent cohorts, while supply-constrained non-employment did not rise. This result points away from a reservation-wage or "pickiness" interpretation of the entry decline. Nevertheless, the LISS sample is small, so the estimates carry considerable noise.

FIGURE B5. Demand-constrained share of youth non-employment across cohorts.

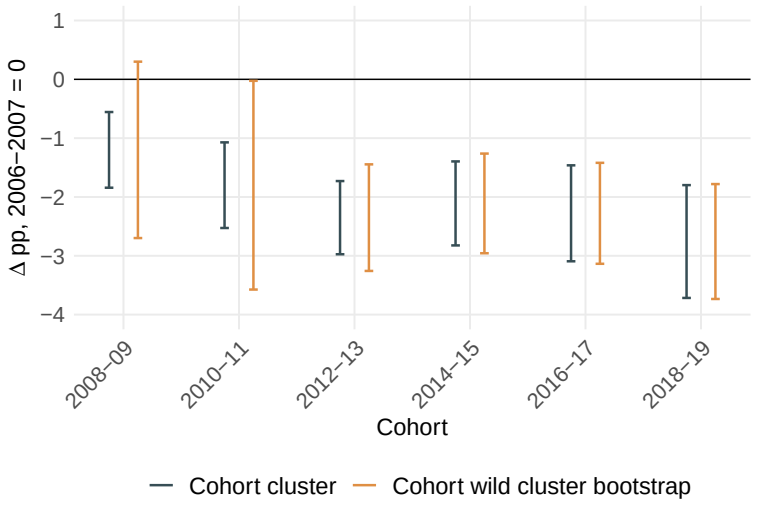


Notes: Estimated cohort effects β_b from Equation (B4), which replaces the employment outcome in Equation (1) with a dummy for demand-constrained non-employment, relative to the 2008 reference, with 95% confidence bands. The data derive from the LISS panel. The sample is non-employed individuals from the 2008 cohort onwards. Number of observations: 1,829; missing 2006 and 2007 cohorts.

B.6. Bootstrapped Clustered Standard Errors

The baseline employment regression in Equation (1) has fourteen cohorts, and therefore fourteen clusters. With few clusters, cluster-robust standard errors may be biased downwards, and tests based on them over-reject (Cameron et al. 2008). As a robustness check, I compute confidence intervals with a wild cluster bootstrap-t at the cohort level, using Webb weights and 10,000 replications (Webb 2023). Figure B6 compares the baseline cohort-clustered confidence intervals with the bootstrap intervals. The bootstrap intervals are wider for the 2008–09 and 2010–11 bins, and the 2008–09 bootstrap interval includes zero. From the 2012–13 bins onwards, the two sets of intervals are similar.

FIGURE B6. Confidence intervals for cohort effects on employment probabilities, by clustering type.



Notes: The figure plots 95% confidence intervals for the estimated cohort employment effects $\hat{\beta}_b$ from Equation (1). The outcome is an indicator for being employed at any point in the year. 'Cohort cluster' intervals use standard errors clustered at the cohort level and critical values from a t-distribution with $14-1=13$ degrees of freedom. 'Wild cluster bootstrap' intervals equal $\hat{\beta}_b \pm c_b^* \cdot \widehat{\text{se}}(\hat{\beta}_b)$, where c_b^* is the 95th percentile of $|t_b^*|$ across 10,000 unrestricted wild cluster bootstrap replications with Webb weights, clustered at the cohort level. Number of observations: 10.6m.

B.7. Selection on Unobservables

The cohort effects in Equation (1) are biased if recent cohorts differ in unobservables such as ability or motivation. Following Oster (2019), I bound the bias from omitting these variables. I replace the cohort-group indicators in Equation (1) with a linear cohort trend C_i because the framework I use is built for continuous variables. The baseline estimates in Figure 1 show no strong non-linearities, which supports the choice.

First, collect the unobservables in W_i. The framework asks how strongly W_i would have to predict cohort membership, relative to the observables X_{it}, to explain the cohort employment effect away. Denote that relative correlation by δ.

Second, estimate a short and a long regression for the linear cohort effect β, with and without the observables X_{it}:

$$\begin{aligned} \text{(short)} \quad & E_{it} = \tilde{\beta}\, C_i + \tilde{u}_{it}, \\ \text{(long)} \quad & E_{it} = \beta\, C_i + X_{it}'\gamma + u_{it}, \end{aligned}$$

in which X_{it} collects the calendar year, education (level and field)-by-age, gender, and migrant-background fixed effects of Equation (1).

Third, bound the bias by comparing the movement in the coefficient on C_i with the movement in the R-squared between the two regressions. Intuitively, if adding X_{it} moves the coefficient but not the R-squared, omitted variable bias is likely to be important. The bias-adjusted cohort effect is:

$$\beta^* \approx \beta - \delta \cdot (\tilde{\beta} - \beta) \cdot \frac{R^2_{\max} - R^2}{R^2 - \tilde{R}^2}, \tag{B5}$$

in which tildes denote the short regression and $R^2_{\max}$ is the R-squared of a hypothetical regression of E_{it} on both X_{it} and W_i. I set $R^2_{\max} = \min\{1, 1.3R^2\}$, following Oster (2019).

Table B1 reports the bias-adjusted estimate β^* for different values of δ, all in percentage points. As δ rises, so that selection on unobservables increasingly follows selection on observables, the employment decline grows from −0.20 to −0.42 pp per cohort. The adjusted effect remains negative even for $\delta = -0.5$. The final column reports the value δ^* that would set β^* to zero. Since $\delta^* = -2.76$ is negative, recent cohorts would have to be positively selected on observables but negatively selected on unobservables almost three times as strongly for the employment decline to disappear.

TABLE B1. Oster (2019) bias-adjusted cohort employment effects.

Correlation: δ	Short model: $\tilde{\beta}$	Long model: β	Bias-adjusted: β^*	δ^*
−0.5	0.05	−0.24	−0.20	−2.76
0	0.05	−0.24	−0.24	−2.76
0.5	0.05	−0.24	−0.29	−2.76
1	0.05	−0.24	−0.33	−2.76
2	0.05	−0.24	−0.42	−2.76

Notes: This table reports the bias-adjusted cohort effect β^* from Equation (B5) for different values of the proportional-selection parameter δ, following Oster (2019). I set $R^2_{\max} = \min\{1, 1.3R^2\}$. $\tilde{\beta}$ is the coefficient from the short regression without controls and β the coefficient from the long regression with the calendar-year, education-by-age, gender, and migrant-background fixed effects of Equation (1). Both regressions use the ramp-adjusted employment indicator (Appendix B.1). δ^* is the value of δ at which the estimated effect equals zero ($\beta^* = 0$). The sample is described in Section 2.1. Number of observations: 10.6 million.

B.8. Robustness of Cohort Employment Effects Specification

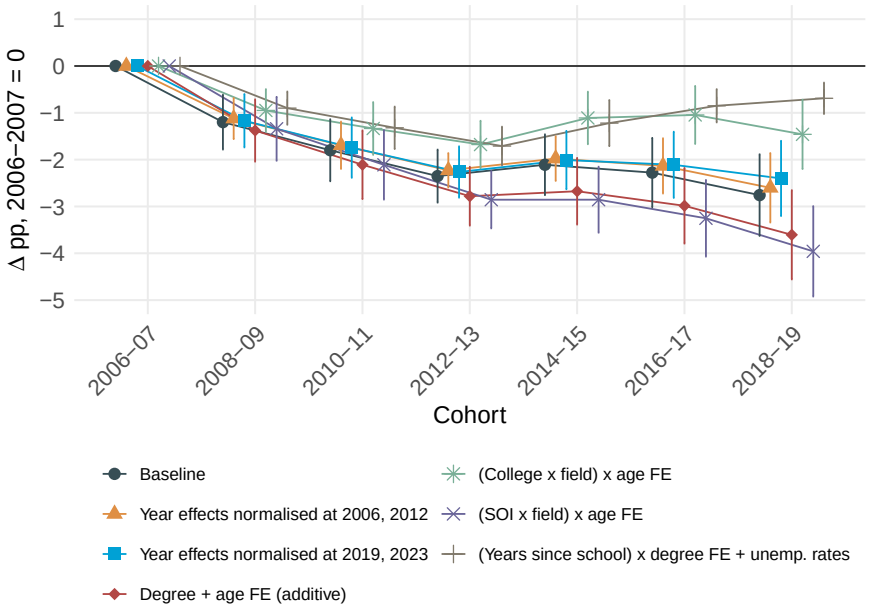
The checks below vary the normalisation of the year effects and the fixed effects. Figure B7 overlays each check, still with 2006–07 as the reference cohorts. 'Baseline' repeats the baseline cohort employment effect estimates.

The cohort effects in Equation (1) are identified from variation in school-leaving age within education-by-age cells. Alternatively, I restrict the year effects to have no linear trend between two chosen years, which assigns the linear time trend to cohorts instead. Figure B7 shows the results when the year effects are normalised at 2006 and 2012 ('Year effects normalised at 2006, 2012') and at 2019 and 2023 ('Year effects normalised at 2019, 2023').[44] The cohort effects are slightly smaller than in the baseline.

Next, I vary the fixed effects in the regression. In 'Degree + age FE (additive)' I replace the education-level-by-field-by-age fixed effects with additive level-by-field and age fixed effects. '(College x field) x age FE' and '(SOI x field) x age FE' replace the education level in the main specification, which distinguishes between 'Low education', 'Medium education' (vocational), 'Bachelor's' and 'Master's', with a college indicator (bachelor's or master's versus lower) or with the SOI level, respectively. The SOI level is Statistics Netherlands' Standaard Onderwijsindeling, which distinguishes across different levels from high school, types of vocational degrees, types of bachelor's and master's degrees, etc. Then, '(Years since schooling) x degree FE + unemp. rates' replaces the degree (level-by-field) by age and the year fixed effects with a degree-by-τ interaction and absorbs the cycle with the aggregate unemployment rate u_t and with the unemployment rate in the year of leaving school, u_c, interacted with $\tau \in \{0, 1, \ldots, 5\}$. The employment decline persists across all specifications, although it is smallest and not monotonic in the unemployment-rate specification.

[44] I choose 2006 as the first year in my sample and 2019 as the last pre-COVID year. I then select the other year for each normalisation as the year with the closest mean employment rate.

FIGURE B7. Robustness of cohort employment effects to regression specification.

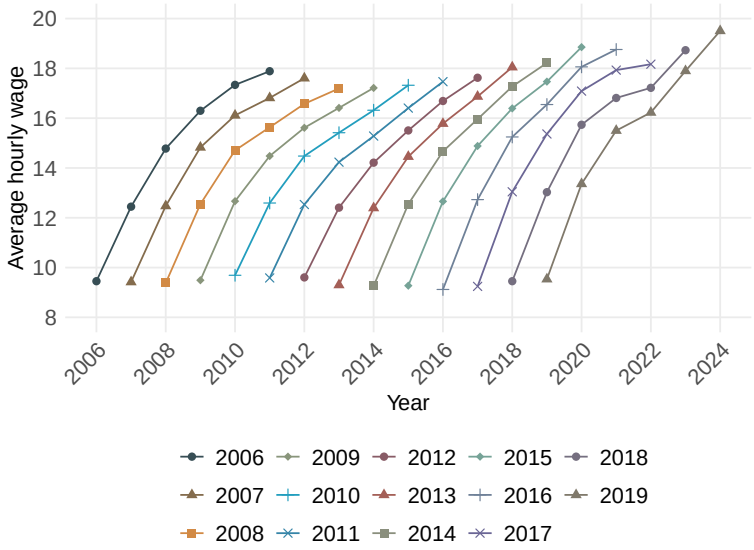


Notes: The figure plots estimated cohort effects $\hat{\beta}_b$ for adjacent two-year cohort bins relative to the 2006–07 reference cohorts, under the baseline specification in Equation (1) and the alternatives described in Appendix B.8: year effects with the linear trend normalised to zero between 2006 and 2012, or between 2019 and 2023, in place of year fixed effects; additive degree and age fixed effects; college-by-field-by-age and SOI-by-field-by-age fixed effects; years-since-schooling-by-degree fixed effects with the aggregate unemployment rate and the unemployment rate in the year of leaving school interacted with years since schooling, in place of age and year fixed effects. All specifications include female and migrant fixed effects. The outcome is an employment indicator, so estimates are in percentage points. The sample is described in Section 2.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 10.6 million in every specification.

B.9. Raw Wage Growth Paths by Cohort

Figure B8 plots mean hourly wages of early-career workers by cohort-year combination $\{c, t\}$. Wages at $\tau = 0$ were flat across cohorts, but as τ increases, the wages of more recent cohorts grew faster.

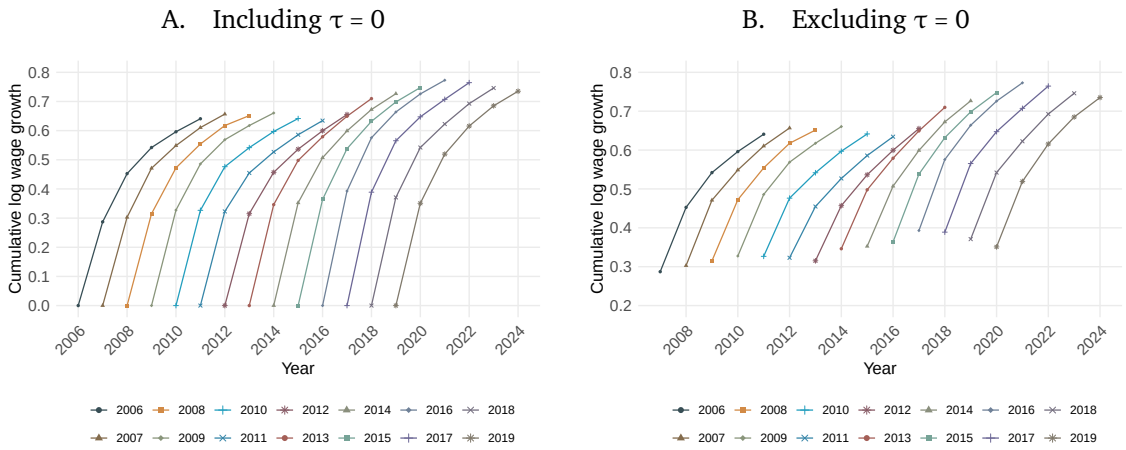
FIGURE B8. Raw hourly wages by cohort and year.



Notes: The figure plots the mean hourly wage, in 2015 euros, for each cohort by calendar year, without controls. The sample is employed individual-years in the early-career panel described in Section 2.1. Number of observations: 10.1 million.

Figure B9 shows the growth more directly. I plot mean cumulative log hourly wage growth since the cohort year by $\{c, t\}$, net of the corresponding growth of age-45 workers to remove aggregate wage movements. Panel A includes the cohort year ($\tau = 0$) and Panel B omits it for readability. Cumulative wage growth rose across cohorts at almost every τ.

FIGURE B9. Raw cumulative growth in log hourly wages by cohort, relative to the cohort year.

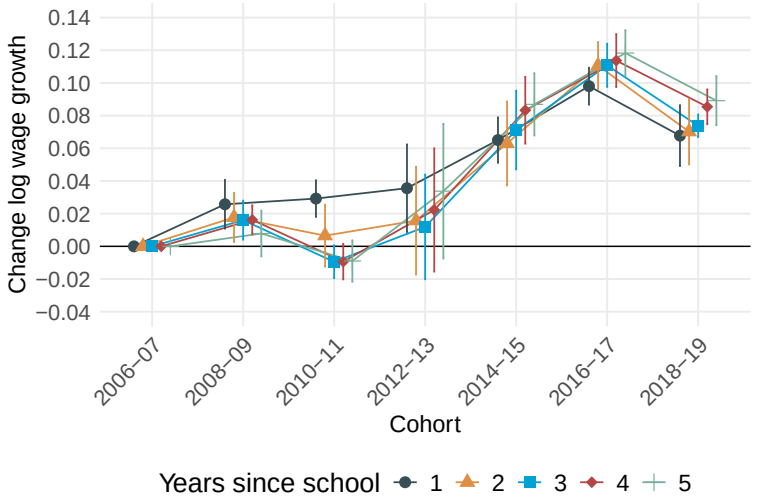


Notes: The figure plots the mean cumulative change in the log hourly wage since the cohort year ($\tau = 0$) for each cohort by calendar year pair. I net out the corresponding change for the age-45 group over the same calendar years. Panel A includes the cohort year $\tau = 0$ and Panel B starts at $\tau = 1$. The sample is employed individual-years in the early-career and age-45 panels described in Section 2.1. Number of observations: 21.4 million.

B.10. Cohort Split

The baseline returns specification splits the fourteen cohorts in half, with $b_0 = \{2006, \ldots, 2012\}$ as the reference and $\{2013, \ldots, 2019\}$ as the comparison group. I do this mainly for ease of presentation. Figure B10 re-estimates Equation (4) with adjacent-cohort bins instead, shown along the x-axis. The returns to experience rose almost monotonically across bins at each τ, and the gap between the $\tau = 1$ and $\tau = 5$ estimates widens with each bin. The rise in early-career returns is therefore not an artefact of the baseline grouping.

FIGURE B10. Change in the returns to early-career experience by adjacent cohort groups.

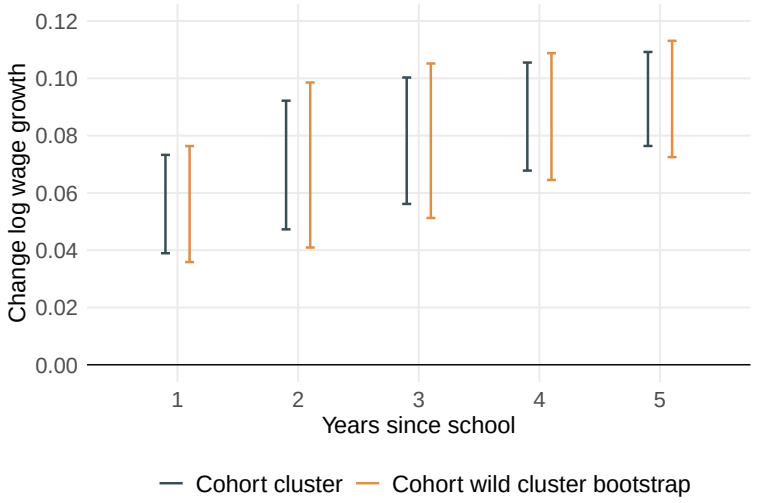


Notes: The figure plots the estimated change in the returns to early-career experience $\hat{\theta}_{b\tau}$ from Equation (4), estimated by ordinary least squares. The estimates are in log points of the hourly wage relative to the 2006–07 reference cohorts. Now I define cohort groups b as adjacent pairs of cohorts (shown on the horizontal axis) rather than the two halves used in the baseline results in Figure 2. Each line corresponds to one value of $\tau \in \{1, \ldots, 5\}$. The sample is described in Section 2.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 21.4 million.

B.11. Bootstrapped Clustered Standard Errors for Returns to Early-Career Experience

The baseline returns to early-career experience regression in Equation (4) has fourteen cohorts, and therefore fourteen clusters. With few clusters, cluster-robust standard errors may be biased downwards, and tests based on them over-reject (Cameron et al. 2008). As a robustness check, I compute confidence intervals with a wild cluster bootstrap-t at the cohort level, using Webb weights and 10,000 replications (Webb 2023). Figure B11 compares the baseline cohort-clustered confidence intervals with the bootstrap intervals. The two sets of intervals are similar at each cohort bin.

FIGURE B11. Confidence intervals for returns to early-career experience, by clustering type.

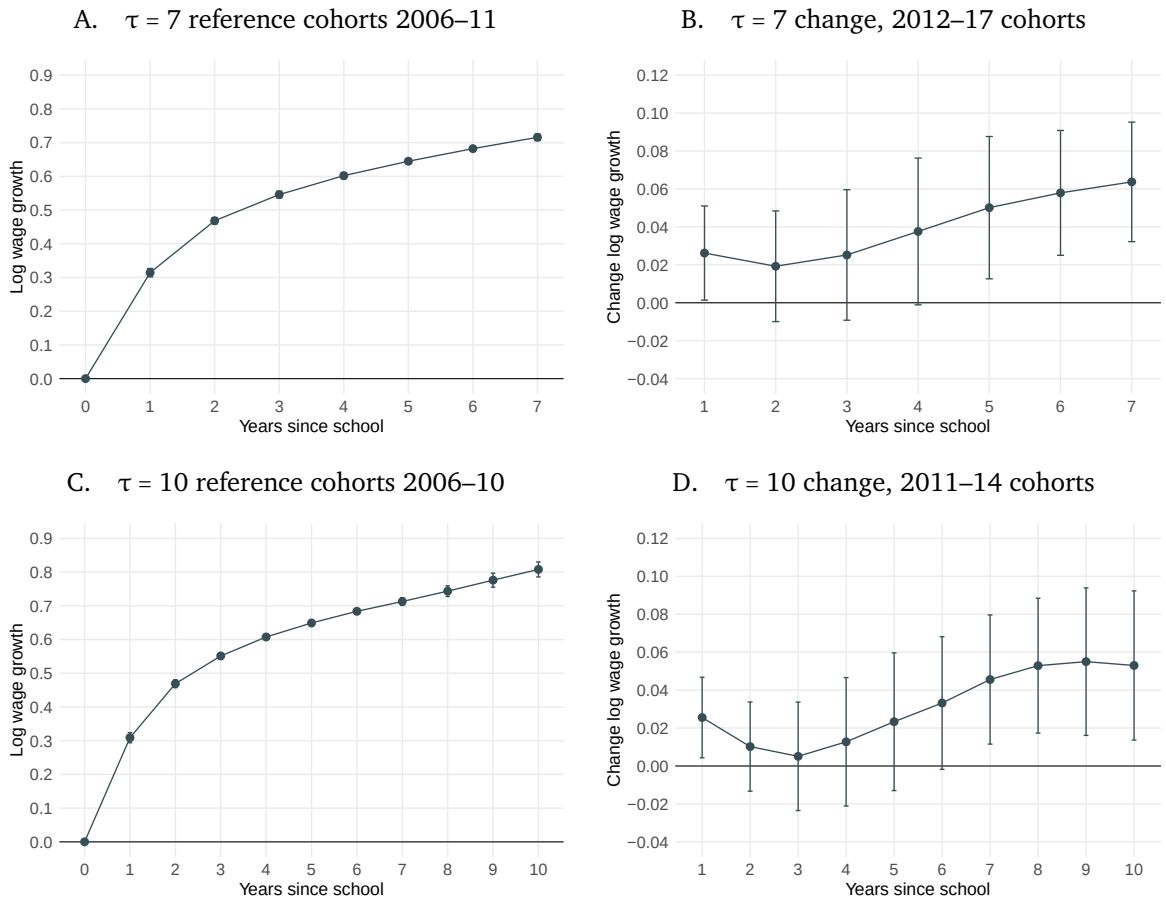


Notes: Estimated 95% confidence intervals for change in returns $\hat{\theta}_{b\tau}$ from Equation (4), estimated by ordinary least squares. 'Cohort cluster' intervals use standard errors clustered at the cohort level and critical values from a t-distribution with 14 − 1 = 13 degrees of freedom. 'Wild cluster bootstrap' intervals equal $\hat{\theta}_{b\tau} \pm c_b^* \cdot \widehat{\text{se}}(\hat{\theta}_{b\tau})$, where c_b^* is the 95th percentile of $|t_b^*|$ across 10,000 unrestricted wild cluster bootstrap replications with Webb weights, clustered at the cohort level. The outcome variable is the log hourly wage, so the confidence intervals are in log points for the 2013–19 cohorts relative to the reference cohorts, netting out the wage growth of the age-45 comparison group over the same calendar years. The sample is the early-career and age-45 panels described in Section 2.1. Number of observations: 21.4 million.

B.12. Extending Potential Experience

The baseline specification compares the returns to early-career experience up to $\tau = 5$ to balance the number of cohorts against the number of years individuals are observed after schooling. I re-estimate the returns to early-career experience regression in Equation (4) up to $\tau = 7$ and $\tau = 10$, which trades fewer cohorts for a longer early career. The sample is the same as in Section 2.1 but ends with the 2017 cohort for $\tau = 7$ and the 2014 cohort for $\tau = 10$. As in the baseline, I split the cohorts into two groups: the reference cohorts are 2006–11 for $\tau = 7$ and 2006–10 for $\tau = 10$. Figure B12 shows that the change in the returns to early-career experience is qualitatively similar to the baseline with $\tau = 5$.

FIGURE B12. Returns to experience up to $\tau = 7, 10$.

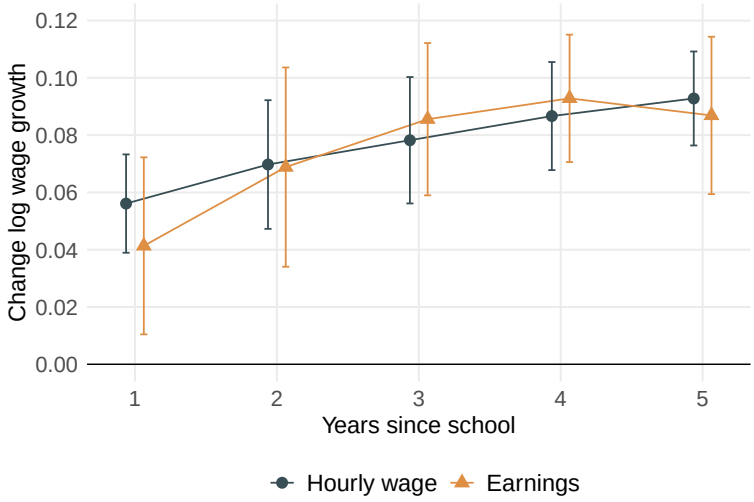


Notes: Estimated change in returns $\hat{\theta}_{b\tau}$ from Equation (4), estimated by ordinary least squares. The outcome variable is the log hourly wage, so each estimate is in log points relative to the reference cohorts. All estimates net out the wage growth of the age-45 comparison group over the same calendar years. I extend potential experience up to $\tau = 7$ years (Panels A and B) and up to $\tau = 10$ years (Panels C and D). The sample is the early-career and age-45 panels described in Section 2.1. I only include up to cohort 2017 for $\tau = 7$ and cohort 2014 for $\tau = 10$. I still split cohorts into two groups: for $\tau = 7$ the reference cohorts are 2006–11 and for $\tau = 10$ the reference cohorts are 2006–10. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 17.6m ($\tau = 7$), 23.9m ($\tau = 10$).

B.13. Wages or Earnings

I re-estimate Equation (4) with log total earnings as the outcome variable, instead of log hourly wages. Figure B13 plots the triple-difference coefficient $\hat{\theta}_{b\tau}$ for earnings next to the baseline wage estimates. The profiles are similar.

FIGURE B13. Change in returns to experience for hourly wages and total earnings.

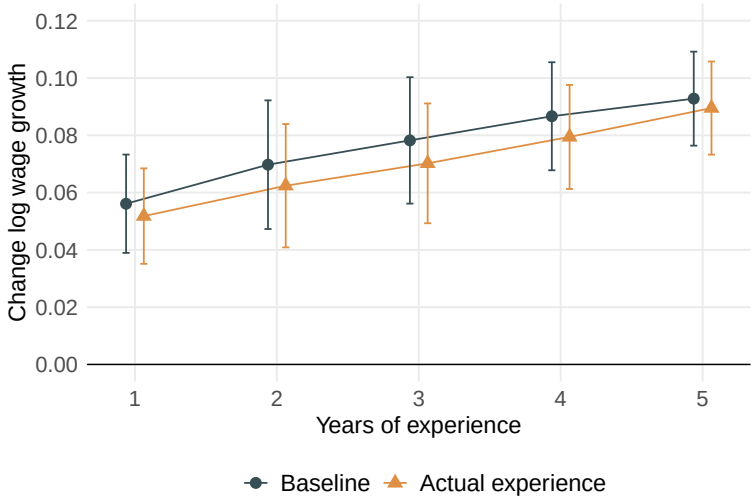


Notes: The figure plots $\hat{\theta}_{b\tau}$ from Equation (4) with the log hourly wage as the outcome ('Baseline') and with log annual earnings as the outcome ('Earnings'), estimated with ordinary least squares. The estimates are log point changes, for the 2013–19 cohorts relative to the 2006–12 reference cohorts. The sample is the early-career and age-45 panels described in Section 2.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 21.4 million.

B.14. Actual Experience

The baseline returns to early-career experience regression in Equation (4) measures experience as years since schooling (i.e., potential experience). Here, I instead regress the log hourly wage on actual experience. Figure B14 shows that the *change* in the actual experience profile ('Actual experience') lies only slightly below the change in the baseline.

FIGURE B14. Change in returns to experience for actual years of experience.

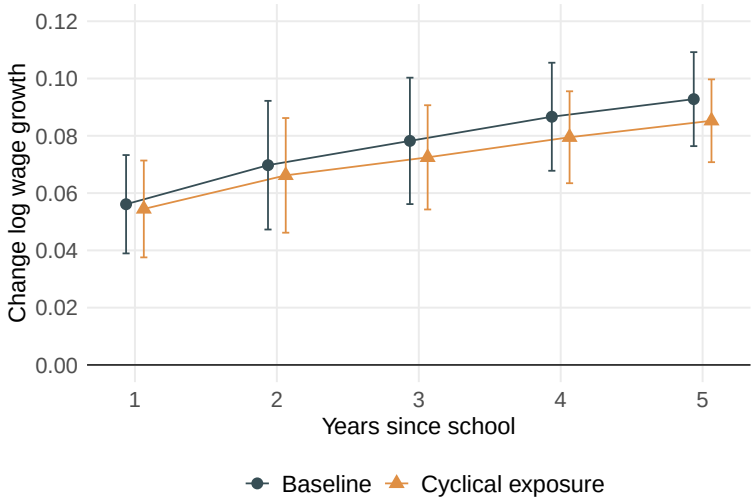


Notes: The figure plots $\hat{\theta}_{b\tau}$ from Equation (4) with experience measured as potential experience ('Baseline') and as actual experience ('Actual experience'), estimated with ordinary least squares. The outcome variable is the log hourly wage, so each estimate measures the change for the 2013–19 cohorts relative to the 2006–12 reference cohorts. The sample is the early-career and age-45 panels described in Section 2.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 21.4 million.

B.15. Differential Cyclical Exposure by Group

The triple-difference design in Equation (4) nets out common business cycle effects through the age-45 control group. However, this specification assumes that the business cycle affects both groups equally. If early-career workers are instead more exposed to aggregate conditions, the control group under-corrects. To proxy for macroeconomic conditions, I add the interaction of the log national unemployment rate with the indicator for an early-career individual I_y to Equation (4) (Borjas 1999; Bratsberg et al. 2006). Figure B15 shows that the resulting 'Cyclical exposure' profile is close to the baseline.

FIGURE B15. Change in returns to early-career experience, with cyclical exposure effects.



Notes: The figure plots $\hat{\theta}_{b\tau}$ from Equation (4), estimated with ordinary least squares. 'Baseline' shows the log hourly wage effects from Equation (4) and 'Cyclical exposure' adds the interaction of the log national unemployment rate and the early-career indicator I_y to Equation (4). The estimates are log point changes for the 2013–19 cohorts relative to the 2006–12 reference cohorts. The sample is the early-career and age-45 panels described in Section 2.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 21.4 million.

B.16. Starting Wages

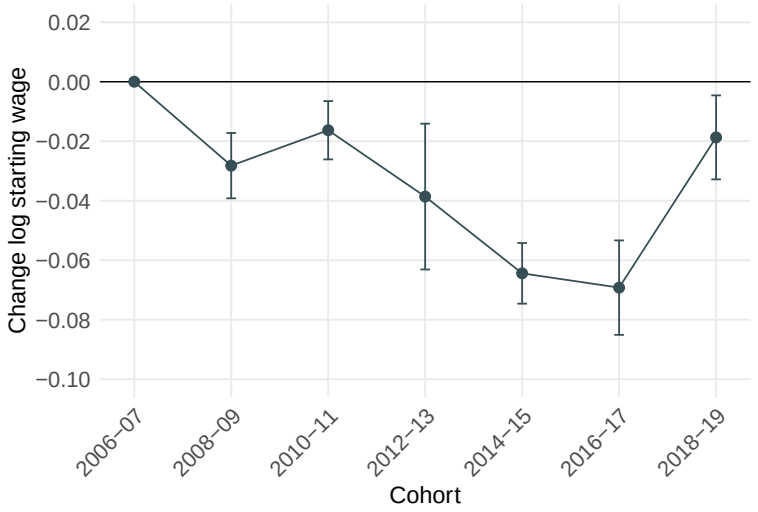
Equation (4) measures each individual's wage growth from the cohort year $\tau = 0$, so the $\tau = 0$ wage is the reference point of every profile. I estimate the cohort profile of the cohort-year wage directly:

$$\ln w_{it} = \sum_{b \in \mathcal{B} \setminus b_0} \kappa_b \left(I_y \times I_b \right) + \phi_c + \upsilon_{ye} + \zeta_g + \nu_m + u_{it}, \tag{B6}$$

in which I_y equals one for early-career workers and zero for the age-45 control group, and I_b indicates the two-year cohort group, with 2006–07 as the reference. Each individual contributes one observation: the cohort year ($\tau = 0$) for early-career workers and the year they turn 45 for the control group. The cohort effects ϕ_c are common to both groups, so κ_b measures the additional change in the starting wage relative to the reference group for early-career workers compared to the age-45 group. The remaining terms absorb composition: υ_{ye} are education level-by-field-by-early-career effects, and ζ_g and ν_m are the gender and migrant-background effects of Equation (1).

Figure B16 plots $\hat{\kappa}_b$. Starting wages declined until the 2016–17 cohort group and ended at around 2 log points below the 2006–07 group for the 2018–19 cohorts. The increase at 2018–19 may reflect the lowering of the full-rate minimum-wage age introduced in 2017–19 (van Bezooijen et al. 2024). Nevertheless, mean reversion from a lower starting wage likely accounts for some of the steeper profile: a back-of-the-envelope calculation suggests that around 20% of the rise can be explained for the 2018–19 group.

FIGURE B16. Cohort effects on hourly wages in the cohort year.

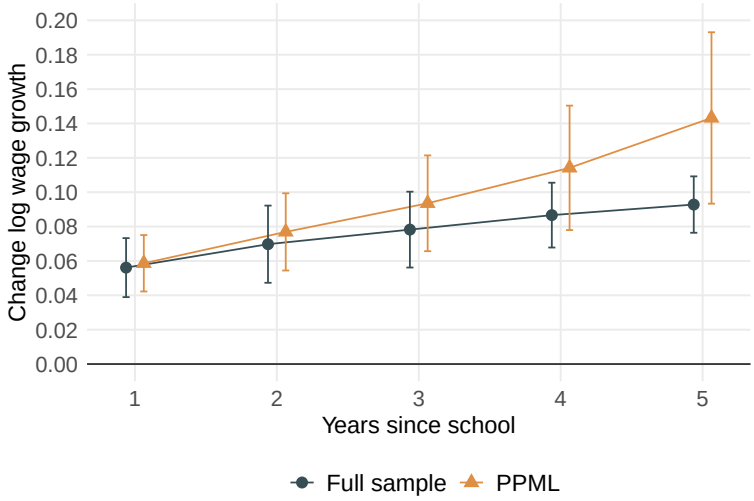


Notes: The figure plots the estimated cohort effects $\hat{\kappa}_b$ on the log hourly wage in the cohort year ($\tau = 0$) from Equation (B6), estimated by ordinary least squares. The estimates are in log points relative to the 2006–07 reference cohorts and net of the corresponding change for the age-45 group. Each individual contributes one observation in their cohort year ($\tau = 0$). The sample is the early-career and age-45 panels described in Section 2.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 2.9 million.

B.17. Functional Form and Selection into Employment

Functional form and zeros. Wages enter Equation (4) in logs, so the steepening could reflect the log transformation rather than actual wage growth. Re-estimating the specification by Poisson pseudo maximum likelihood (PPML) on the hourly wage results in a larger rise in the returns to early-career experience than the baseline at every $\tau \geq 2$ (Figure B17).

FIGURE B17. Change in returns to early-career experience, Poisson pseudo maximum likelihood.



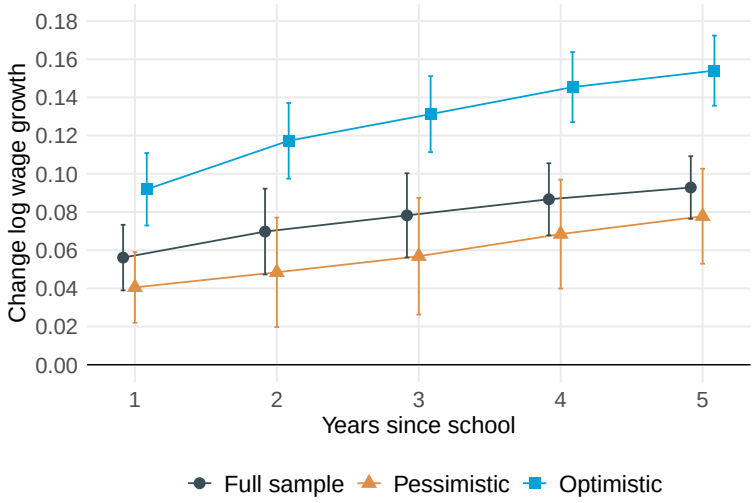
Notes: The figure plots $\hat{\theta}_{b\tau}$ from Equation (4) estimated by ordinary least squares on the log hourly wage ('Baseline') and by Poisson pseudo maximum likelihood on the hourly wage in levels ('PPML'). Both sets of estimates are in log points for the 2013–19 cohorts relative to the 2006–12 reference cohorts. The Poisson sample includes individual-years with zero wages, which the log specification drops. The sample is the early-career and age-45 panels described in Section 2.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 21.4 million (ordinary least squares) and 23.1 million (Poisson).

Selection into employment. Equation (4) is identified from the wage growth of employed individuals. Even though individual fixed effects difference out selection on wage levels for the estimation, here I follow Horowitz and Manski (2000) to impute the counterfactual wage of the non-employed to show that selection into employment does not bias the estimated returns to early-career experience.

In particular, let $[\underline{w}, \overline{w}]_{\tau gb}$ denote the support of observed log hourly wages at potential experience τ for group g and cohorts b. The term g denotes the early-career or age-45 group. For each non-employed individual in the corresponding group and cohorts, I impute their log hourly wage from the observed support. I do this once for the upper bound and once for the lower bound before re-estimating Equation (4) with each imputation. Figure B18 plots the baseline next to the pessimistic ($w^N = \overline{w}$) and optimistic ($w^N = \underline{w}$) imputations. The rise in the

returns to early-career experience stays positive at all $\tau \geq 1$ for both imputations. The bounding exercise suggests that falling employment probabilities across cohorts do not generate the rising returns to experience through selection into employment.

FIGURE B18. Returns to early-career experience under worst-case selection bounds.



Notes: The figure plots $\hat{\theta}_{b\tau}$ from Equation (4) for the 2013–19 cohorts relative to the 2006–12 reference cohorts. I estimate Equation (4) on employed individual-years only ('Full sample') and with the wages of non-employed individual-years imputed at the bounds of the observed wage support following Horowitz and Manski (2000): the upper bound $\overline{w}$ ('Pessimistic') and the lower bound $\underline{w}$ ('Optimistic'). Bounds are computed separately by τ and for the early-career and age-45 groups. The sample is the early-career and age-45 panels described in Section 2.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 21.4 million ('Full sample') and 24.3 million (both bounds).

B.18. Composition of Wage Growth

This subsection shows that (i) the steepening of the returns to experience occurred within firms rather than through reallocation across firms, and (ii) changing worker composition, which could arise if falling employment leaves a pool of employed workers positively selected on wage growth as in Elsby and Shapiro (2012), does not account for the returns steepening.

Within-firm steepening, not reallocation. To separate within-firm wage growth from sorting into high-growth firms, I map residualised wage growth into the decomposition of Foster et al. (2001).

In the first step, I construct a residual wage by subtracting the individual fixed effect from the log wage, $\tilde{w}_{it} \equiv \ln w_{it} - \hat{\alpha}_i$, and average the residuals within each firm j, cohort group b, and potential experience τ to obtain $r^b_{j\tau}$. Aggregating across firms with hours-worked weights $\omega^b_{j\tau}$ gives:

$$R^b_\tau \equiv \sum_j \omega^b_{j\tau} r^b_{j\tau}, \quad \tau \in \{0, \ldots, 5\}, \;\; b \in \mathcal{B}. \tag{B7}$$

The difference $R^b_\tau - R^b_0$ is a residualised analogue of the returns to early-career experience for a cohort group b in Equation (3). Because each term is a weighted average, the difference can be split into a part driven by firm-level residuals (within firm) and a part driven by changes in the weights (across firms):

$$\begin{aligned} R^b_\tau - R^b_0 = {} & \underbrace{\sum_{j \in CO} \omega^b_{j0} \Delta r^b_{j\tau}}_{\text{within}} + \underbrace{\sum_{j \in CO} \left(r^b_{j0} - R^b_0\right) \Delta \omega^b_{j\tau}}_{\text{between}} + \underbrace{\sum_{j \in CO} \Delta \omega^b_{j\tau} \Delta r^b_{j\tau}}_{\text{cross}} \\ & + \underbrace{\sum_{j \in EN} \omega^b_{j\tau} \left(r^b_{j\tau} - R^b_0\right) + \sum_{j \in EX} \omega^b_{j0} \left(r^b_{j0} - R^b_0\right)}_{\text{net entry}}, \end{aligned} \tag{B8}$$

in which CO, EN, and EX denote firms that continue, enter, and exit, respectively, between $\tau = 0$ and some $\tau > 0$.

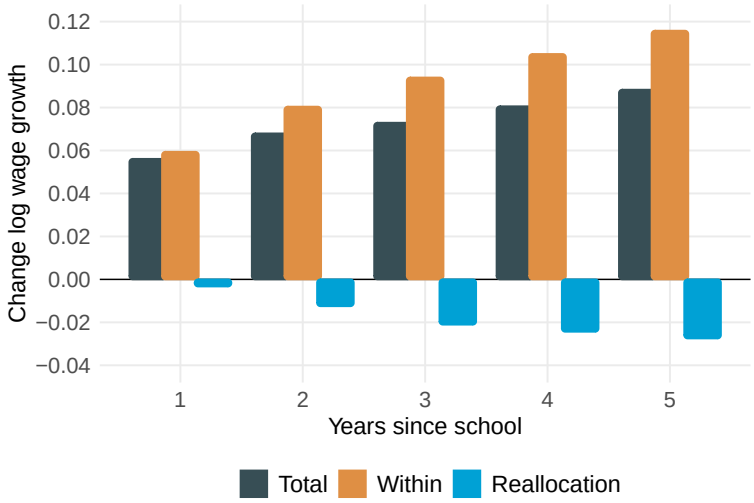
In the second step, I difference each term in Equation (B8) between the later (b) and earlier (b_0) cohort groups: $\Delta_\tau \equiv (R^b_\tau - R^b_0) - (R^{b_0}_\tau - R^{b_0}_0)$. The change in the returns to experience across cohorts, Δ_τ, then decomposes additively into a within-firm component and a reallocation component (the sum of the between, cross, and net-entry terms). The within-firm component captures steeper wage growth inside continuing firms; the reallocation component captures recent cohorts sorting into high-growth firms at a higher rate than earlier cohorts.

In the final step, I compute the decomposition separately for early-career and age-45 workers and subtract each age-45 term from the corresponding early-career term. This step

nets out economy-wide trends in mobility and wage growth in the spirit of the triple-difference design in Section 2.3.

Figure B19 plots Δ_τ and the contribution of each component for $\tau \in \{1, \ldots, 5\}$. The residualised Δ_τ closely tracks the regression estimates in Figure 2, which validates the mapping. The within-firm component accounts for more than 100% of the rise in returns at every τ. Reallocation offsets part of the within-firm contribution. Had recent cohorts increasingly matched into high-growth firms, the gain would have loaded onto the reallocation component. The decomposition points to a shift in how employers reward early-career experience, rather than a change in matching between early-career workers and employers, as the main driver of the rise.[45]

FIGURE B19. Decomposition of the cross-cohort change in returns to experience.



Notes: The figure plots the change across cohort groups in the residualised returns to early-career experience, Δ_τ, for $\tau \in \{1, \ldots, 5\}$, and its decomposition from Equation (B8). Each term is differenced between the 2013–19 and 2006–12 cohort groups and netted against the same term computed for the age-45 group. 'Within' is the within-firm contribution of continuing firms, and 'Reallocation' is the sum of the between, cross, and net-entry terms from differencing Equation (B8) across cohort groups. Wages are residualised on individual fixed effects before aggregation, and firms are weighted by their share of hours worked. The sample is the early-career and age-45 panels described in Section 2.1. Number of observations: 21.4 million.

An AKM decomposition of wage growth. To separate individual and firm contributions to the change in wage returns to early-career experience, I estimate an Abowd et al. (1999) model on wage growth instead of wage levels. The outcome is the log wage change $\Delta \ln w_{it}$ between $t-1$ and t, which I regress on a worker fixed effect α_i and a fixed effect ψ_{jb} for the

[45] A large literature emphasises reallocation, in particular mobility across firms, as a key source of early-career wage growth (Topel and Ward 1992; Bagger et al. 2014). Across cohorts, I find that mobility contributes negatively to the change in the returns to early-career experience. This result is in line with Baksy et al. (2026), who document declining mobility and wage gains from mobility across firms in the US.

firm j by cohort group b:

$$\Delta \ln w_{it} = \alpha_i + \psi_{jb} + \beta \tau_{it}^2 + r_{it}. \tag{B9}$$

Importantly, the firm-cohort-group fixed effect is defined on the cohort group of the worker i. The quadratic in potential experience τ_{it} accounts for concave early-career wage growth, which is needed because the decomposition applies to the linear component of wage growth.

The worker effect α_i captures individual wage growth that does not depend on the employer, and the firm effect ψ_{jb} the wage growth that firm j offers to cohort group b. Firm effects are identified from firm-to-firm movers in the largest connected set. Because each worker belongs to one cohort group, I normalise the firm effects so that the mean effect in the accommodation and food services sector is zero for each group b, following Card et al. (2016).[46] As elsewhere in the paper, I estimate Equation (B9) for the 2006–12 and 2013–19 cohort groups separately.

Using the estimated fixed effects from Equation (B9), I average Equation (B9) within each cohort group b and difference across them:

$$\Delta\left(\overline{\Delta \ln w_{it}}\right) = \Delta\left(\overline{\alpha_i}\right) + \Delta\left(\overline{\psi_{jb}}\right) + \Delta\left(\overline{\beta \tau_{it}^2 + r_{it}}\right), \tag{B10}$$

in which the overline denotes a mean within b, and the $\Delta()$ denotes a difference across b. The left-hand side of Equation (B10) is the change in mean early-career wage growth. The worker component, $\Delta\left(\overline{\alpha}\right)$, captures changes in worker composition. The firm component, $\Delta\left(\overline{\psi_{jb}}\right)$, captures the change in the annual wage growth that firms offer to early-career workers. Table B2 shows the decomposition for early-career workers, after subtracting the corresponding age-45 components. 'Better' worker composition in recent cohorts can account for 34% of the steeper wage growth, but the firm component dominates and accounts for a much larger share.

[46] The sector is the lowest-paying sector in the data and is conventionally treated as the sector in which firms pay no rents. I extend the assumption to wage growth.

TABLE B2. Decomposition of change in early-career wage growth into worker and firm components.

	Log points	Share of total
Change in mean wage growth	0.009	100%
Worker component: $\Delta\bar{\alpha}$	0.003	34%
Firm-by-cohort-group component: $\Delta\bar{\psi}$	0.013	145%
Within firms	0.016	181%
Between firms	−0.003	−36%
Residual	−0.005	−78%

Notes: The table decomposes the change in mean annual early-career wage growth between the 2006–12 and 2013–19 cohort groups using Equation (B9). Equation (B9) is an extension of the two-way fixed effects model from Abowd et al. (1999) to decompose wage *growth* into worker effects α_i and firm-by-cohort-group effects ψ_{jb}. Cohort groups are the same 2006–12 versus 2013–19 groups as in the baseline exercises. The first column reports each component's contribution in log points (per year) to the change in the mean wage growth across the two cohort groups. The second column shows the share of the total change. The firm component is further split into within-firm changes and between-firm reallocation. Firm effects are normalised to a mean of zero in the accommodation and food services sector within each cohort group, following Card et al. (2016). The model is estimated on the largest connected set of firms, separately for early-career and age-45 workers, and each age-45 component is subtracted from the corresponding early-career component. The sample is the early-career and age-45 panels described in Section 2.1. Number of observations: 12.4m.

B.19. Institutions

Two Dutch institutions could steepen the early-career wage profile: collective bargaining and the spread of temporary employment contracts. In this subsection, I show that neither accounts for the rise in the returns to early-career experience across cohorts.

Collective bargaining. Collective bargaining agreements (CBAs) cover around 70–80% of employees in the Netherlands and are often extended to all workers in a sector. Typically, CBAs specify minimum wage increases rather than wage levels. However, if the content of the agreements had shifted across cohorts, CBAs could explain the rising returns to early-career experience. For example, CBAs may have become more backloaded.

The institutional structure of CBAs delivers a sharp test of this hypothesis. If CBAs were responsible, the change in the returns to early-career experience should be similar for all firms within a sector-level CBA. Thus, the variance in the firm-level returns within sectors should be close to zero. Let $\Delta r_{j\tau}$ denote the change in the mean wage growth from $\tau = 0$ to $\tau > 0$ in firm j between cohort groups b_0 and b. The mean is defined in the same way as in the within-between decomposition in Appendix B.18. Assigning each firm to its sector s from the business registry, the variance of $\Delta r_{j\tau}$ decomposes into within- and between-sector terms:

$$\text{Var}(\Delta r_{j\tau}) = \underbrace{\text{Var}\left(\Delta r_{j\tau} - \mathbb{E}_s[\Delta r_{j\tau}]\right)}_{\text{Within-sector}} + \underbrace{\text{Var}\left(\mathbb{E}_s[\Delta r_{j\tau}]\right)}_{\text{Between-sector}}. \tag{B11}$$

Table B3 reports the decomposition, weighted by hours worked. Between 87 and 93% of the variance occurs *within* sectors (98–99% unweighted). CBA content therefore cannot account for the rise in the returns to early-career experience.

TABLE B3. Within- and between-sector decomposition of change in early-career returns to experience.

	Weighted		
	(1)	(2)	(3)
Years since schooling	Total variance	Within	Between
1	0.115	87%	12%
2	0.104	93%	7%
3	0.102	91%	8%
4	0.103	91%	8%
5	0.106	91%	8%

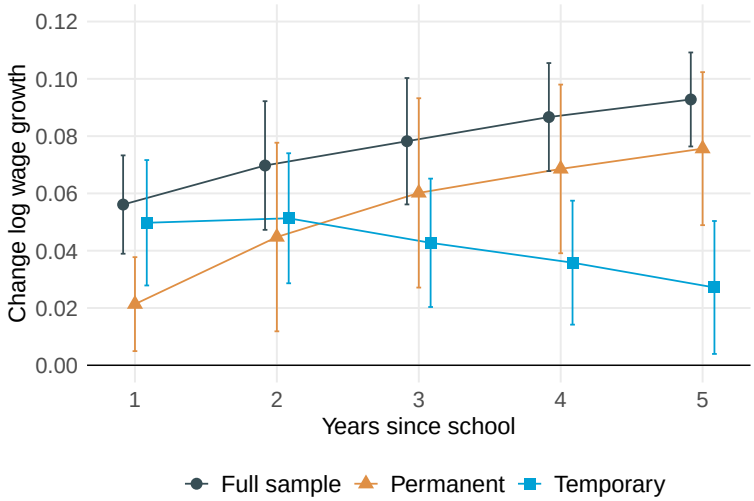
Notes: The table decomposes the variance across firms of $\Delta r_{j\tau}$, the change between the 2006–12 and 2013–19 cohort groups in firm j's mean cumulative *early-career* wage growth from $\tau = 0$ to τ. Column 1 reports the total variance. Columns 2 and 3 show the split into within- and between-sector contributions using Equation (B11), separately for each τ. Firm-level growth is net of the corresponding growth of the firm's age-45 workers. Sectors are two-digit codes of the Dutch Standard Industrial Classification (Standaard Bedrijfsindeling). Results are similar at the four-digit level. Each firm's contribution to the total variance is hours-weighted; unweighted, the within-sector share is 98–99%. The sample is the early-career and age-45 panels described in Section 2.1. Number of observations: 21.4m.

Temporary contracts. The Netherlands, like many OECD economies, has seen a rapid expansion of temporary contracts since the late 1980s, particularly among young workers (Bolhaar et al. 2018). Since wages on a temporary contract may be fixed over its duration, a worker may receive a discrete gain when the contract converts to a permanent one (Daruich et al. 2023; Figueiredo et al. 2026). Legally, a conversion must occur within three years of tenure, which would induce a discrete wage gain during the $\tau \leq 5$ window.[47]

To test the channel, I estimate Equation (4) separately for permanent and temporary contracts. Figure B20 shows a steepening across cohorts within both contract types at every $\tau \in \{1, \ldots, 5\}$. However, the *change* in the temporary-contract profile falls as τ increases, which suggests that permanent contracts account for most of the change in returns to early-career experience that I document. Temporary-to-permanent conversions may still account for the gap between the 'Full sample' and 'Permanent' lines.

[47] The cap was three years before 2015, two years from 2015 to 2019, and three years from 2020 onwards.

FIGURE B20. Change in the returns to experience by contract type.



Notes: The figure plots $\hat{\theta}_{b\tau}$ from Equation (4) estimated separately for individual-years on permanent contracts ('Permanent') and temporary contracts ('Temporary'), alongside the baseline ('Full sample'). The estimates are in log points of the hourly wage for the 2013–19 cohorts relative to the 2006–12 reference cohorts. Contract type is that of the primary employer in the year. The sample is the early-career and age-45 panels described in Section 2.1. Bars show 95% confidence intervals with standard errors clustered at the cohort level. Number of observations: 21.4 million ('Full sample'), 12.6 million ('Permanent'), and 8.2 million ('Temporary').

Appendix C. Digital Technologies Evidence

This appendix reports additional results and robustness checks for the firm-level (Section 3) and the worker-level (Section 4) DiD design and results.

C.1. Treated and Control Firms

The stacked DiD compares treated firms with not-yet-treated firms observed in the same calendar window. Identification requires parallel pre-event trends. Not-yet-treated firms are a closer comparison than never-treated firms. To show this, I follow Bessen et al. (2025) and estimate:

$$\Delta \ln y_{jt} = \delta_{st} + \kappa I_{\text{ever spike}_j} + u_{jt}, \tag{C1}$$

in which $\Delta \ln y_{jt}$ is the year-to-year change in the log of an outcome such as revenue, employment, or value added for firm j at time t. The term δ_{st} stands for a two-digit sector-by-year fixed effect, and $I_{\text{ever spike}_j}$ denotes a dummy equal to one if firm j ever experiences an investment spike, and zero otherwise.

The estimates for κ are shown in Table C1. Firms that ever experience a spike grow faster in terms of revenue and value added, but not employment. Table 1 in the main text compares mean characteristics of treated and not-yet-treated firms at event time $k = -1$ and shows the two groups are more comparable in levels. The pre-trends in Figure 4 show similar *growth* paths.

TABLE C1. Ever-treated versus never-treated firms.

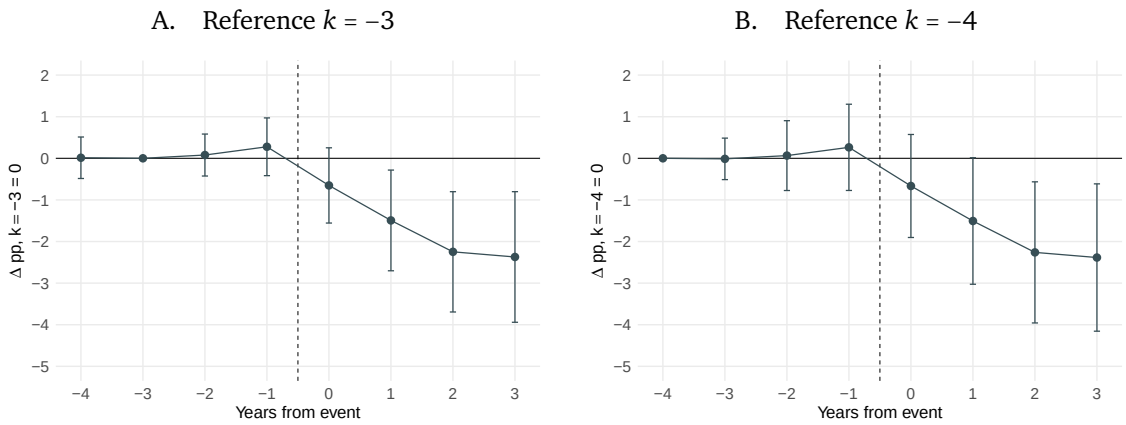
	(1) $\Delta \ln$(Revenue)	(2) $\Delta \ln$(Employment)	(3) $\Delta \ln$(Value Added)
Ever spike	0.0185***	−0.0005	0.0164***
	(0.0010)	(0.0009)	(0.0016)
Sector-year FE	Yes	Yes	Yes
Observations	676k	773k	528k

Notes: The table reports ordinary least squares estimates of κ from Equation (C1). The outcome variable is the year-to-year change in log revenue (Column 1), log employment (Column 2), and log value added (Column 3). I compare firms that ever have a digital-technologies investment spike, as defined in Equation (5), and firms that never do. I add two-digit sector-by-year fixed effects. The sample is all firms in the Production Statistics panel with at least three years of investment data, from 2000 to 2023. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

C.2. Alternative Reference Periods

The dynamic DiD in Figure 4 uses $k = -1$ as the reference year. Identification requires no anticipatory adjustment in the EC share. Anticipatory adjustment would appear at $k = -1$, which is the omitted period in Equation (7). Re-normalising the dynamic specification to use $k = -3$ or $k = -4$ instead of $k = -1$ as the reference period shows that $\hat{\delta}_{k=-1}$ is small and not statistically different from zero. Figure C1 shows that in both cases, the coefficients at $k = -2$ and $k = -1$ are small and statistically insignificant (0.28 and 0.26 pp at $k = -1$). Moreover, the post-event path is unchanged relative to the baseline.

FIGURE C1. Early-career employment share around investment spikes, alternative reference years.

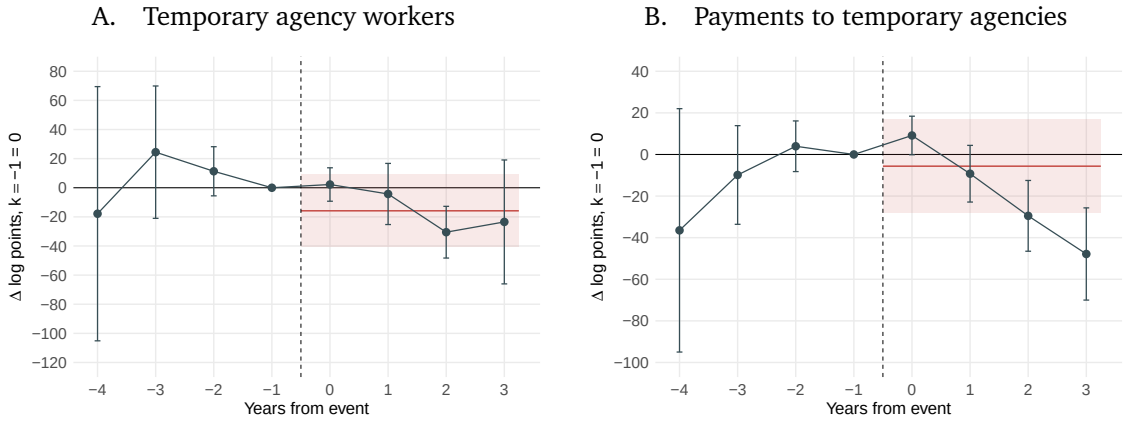


Notes: The figure plots the dynamic difference-in-differences coefficients $\hat{\delta}_k$ from Equation (7) by event time $k = t - s$. Each estimate is relative to the omitted year $k = -3$ (Panel A) and $k = -4$ (Panel B) instead of $k = -1$. The outcome is the early-career employment share in percentage points (0–100). The sample is the stacked panel described in Section 3.2, with event years s from 2010 to 2019. Observations are weighted by the firm's early-career employment at $k = -1$. Bars show 95% confidence intervals with standard errors clustered at the firm level. Number of observations: 263k.

C.3. Outsourcing and Temporary Agency Workers

The DT investment event may lead firms to move employment off the payroll through outsourcing, in which case the fall in the EC share would be a relabelling of the payroll rather than a change in labour demand. If so, the number of temporary agency workers, and payments to staffing agencies, should rise. Figure C2 re-estimates the pooled and dynamic DiDs (Equations (6) and (7)) with the log number of temporary agency workers (Panel A) and log payments to temporary agencies (Panel B) as outcomes. The pooled estimates are negative and statistically insignificant, which suggests that firms do not replace early-career employees with temporary agency workers after the event.

FIGURE C2. Temporary agency workers and payments to staffing agencies around digital-technologies investment spikes.



Notes: The figure plots the dynamic difference-in-differences coefficients $\hat{\delta}_k$ from Equation (7), by event time $k = t - s$ and relative to $k = -1$. I overlay the pooled estimate $\hat{\theta}$ from Equation (6) on the post-event window with its 95% confidence interval. The outcome is the log number of temporary agency workers (Panel A) and log payments in thousands of euros to staffing agencies (Panel B), each multiplied by 100 so that estimates are in log points. The sample is the stacked panel described in Section 3.2, with event years s from 2010 to 2019, restricted to firm-years reporting the outcome. Observations are weighted by the firm's early-career employment at $k = -1$. Bars show 95% confidence intervals with standard errors clustered at the firm level. Number of observations: 120k (Panel A) and 124k (Panel B).

C.4. Mechanisms and Heterogeneity

To probe the mechanism behind the baseline DiD estimate, I estimate the baseline DiD for various other outcomes. Columns 2 and 3 of Table C2 report the change in the log employment of early-career workers and of all remaining workers. After the event, early-career employment falls by around 11 log points, while remaining employment rises by 2 log points. The latter estimate is not statistically significant at the 10% level. Thus, the decline in the EC share after an investment spike reflects fewer early-career workers rather than more remaining workers.[48] The pooled estimate for total employment is 0.1 log points (SE 5.9). Column 4 shows that the decline in the early-career share of hours worked is similar to the baseline EC share after a spike.

TABLE C2. Difference-in-differences estimates for employment in levels and hours worked.

	(1) Baseline	(2) Early-career employment	(3) Remaining employment	(4) Hours worked
Treated	−1.64***	−11.68*	2.38	−1.62***
× Post	(0.57)	(6.59)	(6.09)	(0.56)
Units	pp	log points	log points	pp
Fixed effects				
Firm-event-year	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Observations	263k	242k	263k	263k

Notes: The table reports the pooled difference-in-differences estimate $\hat{\theta}$ from Equation (6). The outcome is the early-career employment share in percentage points (Column 1), log early-career employment (Column 2), log employment of all other workers (Column 3), and the early-career share of hours worked in percentage points (Column 4). Each outcome is multiplied by 100 so that estimates are in percentage points or log points. The sample is the stacked panel described in Section 3.2, with event years s from 2010 to 2019. Observations are weighted by the firm's early-career employment at $k = -1$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Two splits of the baseline regression further probe whether the decline in the EC share is (i) specific to a period, or (ii) driven by financial pressure from the investment outlay.

Column 2 of Table C3 interacts the treatment indicator with an indicator for event years

[48] For comparison, Bessen et al. (2025) find that total employment falls by around 20% over the four years after an event in firms with fewer than 500 workers, with no significant change in larger firms. My sample period and outcomes differ from theirs, but the results add that the incidence falls on early-career workers in particular.

from 2013 onwards. The interaction estimate is −0.29 pp with a standard error of 1.12, so the effect is not distinguishable across the two periods. The pre-2013 effect becomes imprecise after the split, although its magnitude is somewhat comparable to the baseline in Column 1.

Column 3 asks whether the decline in the EC share reflects financial constraints following the investment outlay. If so, the decline should be larger in firms with less cash on hand. I split the sample at the median of accumulated profit over event time $k \in \{-4, \ldots, -1\}$, pooling all observations in the stacked sample. Firms below the median show an EC share decline of 1.97 pp, and the differential for above-median firms is 0.41 pp with a standard error of 0.77. I therefore cannot reject that above- and below-median firms reduce their EC share by the same amount. Even if the differential were significant, the effect for above-median firms would remain below zero.

TABLE C3. Difference-in-differences estimates by period and by pre-event profitability.

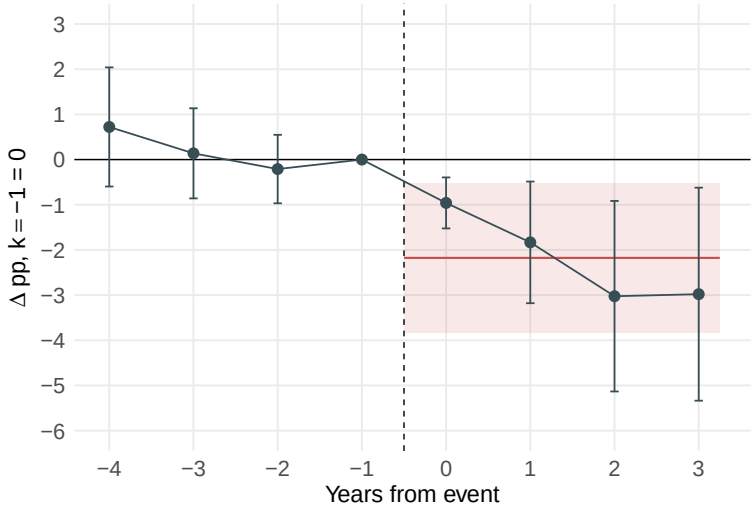
	(1) Baseline	(2) Before and after 2013	(3) Pre-event profit
Treated	−1.64***	−1.39	−1.97***
× Post	(0.57)	(0.94)	(0.51)
Post-2012 /		−0.29	0.41
Above-median		(1.12)	(0.77)
Units	pp	pp	pp
Fixed effects			
Firm-event-year	Yes	Yes	Yes
Year	Yes	Yes	Yes
Observations	263k	263k	263k

Notes: Column 1 repeats the baseline difference-in-differences estimate $\hat{\theta}$ from Equation (6). Column 2 adds the interaction of the treatment-by-post indicator with an indicator for event years $s \geq 2013$. In Column 2, the first row is the effect for event years before 2013 and the second row the differential for later event years. Column 3 adds the interaction with an indicator for firms whose accumulated operating profit over event time $k \in \{-4, \ldots, -1\}$ lies above the median of the pooled stacked sample. In Column 3, the first row is the effect for below-median firms and the second row the differential for above-median firms. The outcome is the early-career employment share in percentage points (0–100). The sample is the stacked panel described in Section 3.2, with event years s from 2010 to 2019. Observations are weighted by the firm's early-career employment at $k = -1$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

C.5. Digital Technologies Spiking Firms Only

Here I restrict the sample from the baseline investment spike results to firms without a spike in materials expenditure, operating profit, tangible capital investment, or outsourcing. The pooled DiD estimate for this subsample is -2.18^{***} pp. Figure C3 shows that the dynamic path resembles the baseline effects in Figure 4.

FIGURE C3. Early-career employment around investment spike, excluding firms with other events.



Notes: The figure plots the dynamic difference-in-differences coefficients $\hat{\delta}_k$ from Equation (7) by event time $k = t - s$, relative to $k = -1$. The sample includes firms from the baseline EC share results restricted to those firms without a spike in materials expenditure, operating profit, tangible capital investment, or outsourcing expenditure. The pooled estimate $\hat{\theta}$ from Equation (6) for this sample is overlaid on the post-event window with its 95% confidence interval. The outcome is the early-career employment share in percentage points. Observations are weighted by the firm's early-career employment at $k = -1$. Bars show 95% confidence intervals with standard errors clustered at the firm level. Number of observations: 107k.

C.6. Econometric Robustness

Table C4 varies the event definition. Column 1 repeats the baseline estimate for the decline in the EC share. For Column 2, I randomly assign each treated firm an event year within its pre-event period. The placebo estimate is small and insignificant. Column 3 defines the spike with a threshold of two times the leave-one-out mean, rather than three times, in Equation (5). Column 4 defines treatment as the first year in which a firm's DT investment exceeds the 80th percentile of the pooled DT investment distribution. Column 5 defines the event using investment per worker. All definitions generate a decline in the EC share consistent with the baseline, but with varying size and precision.

Table C5 varies the regression specification and the sample. Column 1 replaces year fixed effects with year-by-event-year fixed effects, Column 2 is an unweighted estimate, Column 3 uses a balanced panel of firms to address attrition, and Column 4 restricts the sample to firms that always employ at least 50 workers. Column 5 replaces the stacked estimator with the Callaway and Sant'Anna (2021) estimator. Every check shows a decline in the EC share, mostly comparable to the baseline.

TABLE C4. Robustness of the early-career employment share estimate: placebo and alternative event definitions.

	(1) Baseline	(2) Placebo	(3) $R = 2$	(4) 80th percentile	(5) Per worker
Treated	−1.64***	−0.18	−0.83**	−0.93***	−1.00
× Post	(0.57)	(0.36)	(0.39)	(0.34)	(0.70)
Units	pp	pp	pp	pp	pp
Fixed effects					
Firm-event-year	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes
Observations	263k	97k	333k	120k	183k

Notes: The table reports the pooled difference-in-differences estimate $\hat{\theta}$ from Equation (6) for various robustness checks. Column 1 repeats the baseline. Column 2 is a placebo that assigns each treated firm a randomly drawn event year within its pre-event window. Column 3 lowers the spike threshold in Equation (5) from three to two times the firm's leave-one-out mean. Column 4 defines the event as the first year in which a firm's digital-technologies investment exceeds the 80th percentile of the pooled investment distribution. Column 5 defines spikes using investment per worker rather than investment. The outcome in every column is the early-career employment share in percentage points (0–100). Each column re-stacks the panel of treated and not-yet-treated firms as described in Section 3.2, with event years s from 2010 to 2019, so the number of observations varies with the event definition. Observations are weighted by the firm's early-career employment at $k = -1$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

TABLE C5. Robustness of the early-career employment share estimate: fixed effects, weights, sample, and estimator.

	(1) Year-by-event-year	(2) Unweighted	(3) Balanced panel	(4) 50+ workers	(5) Callaway and Sant'Anna
Treated	−1.87***	−0.60***	−1.67**	−1.73**	−3.13***
× Post	(0.61)	(0.09)	(0.76)	(0.70)	(0.67)
Units	pp	pp	pp	pp	pp
Fixed effects					
Firm-event-year	No	Yes	Yes	Yes	
Year	Yes	Yes	Yes	Yes	
Year-event-year	Yes	No	No	No	
Observations	263k	417k	147k	103k	42k

Notes: The table reports the pooled difference-in-differences estimate $\hat{\theta}$ from Equation (6) under alternative fixed effects, weights, samples, and estimators. Column 1 replaces the year fixed effects with year-by-event-year fixed effects. Column 2 shows unweighted estimates. Column 3 restricts the sample to a balanced panel in event time. Column 4 restricts the sample to firms with at least 50 workers in every year. Column 5 replaces the stacked design with the Callaway and Sant'Anna (2021) estimator, using not-yet-treated firms as controls and aggregating group-time effects into a single average. The outcome is the early-career employment share in percentage points (0–100). The sample is the stacked panel described in Section 3.2, with event years s from 2010 to 2019. The Callaway and Sant'Anna estimator uses a non-stacked panel, but otherwise the same event definition. Unless stated, observations are weighted by the firm's early-career employment at $k = -1$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

C.7. Local Spillovers

The stable unit treatment value assumption (SUTVA) in the DiD design fails if spikes at treated firms change the EC share of control firms. Two local channels could do so. Under *absorption*, early-career workers who are no longer hired by a treated firm find jobs at nearby not-yet-treated firms, raising the control group's EC share and overstating the negative treatment effect. Under *contagion*, nearby firms respond to a competitor's spike by investing themselves, so the control group is partially treated and the effect is understated. I check for both channels in two ways.

First, I ask whether controlling for the local intensity of spikes changes the baseline estimate. Let S_{mt} denote the share of firms in municipality-sector m that have a spike in year t.[49] If a firm's EC share is unaffected by spikes in similar firms elsewhere in the municipality, adding S_{mt} for each treated and control firm j should leave the treatment estimate unchanged. Column 2 of Table C6 shows an estimate of −1.82 pp, close to the baseline in Column 1. Local spike intensity absorbs none of the treatment effect.

Second, I test directly whether control firms respond to nearby spikes. Columns 3 and 4 regress the EC share of not-yet-treated firms on S_{mt} and three of its lags, with the same fixed effects as the baseline. If non-hired early-career workers were absorbed by nearby firms, these coefficients would be positive. All coefficients are small, statistically insignificant, and negative. Local spillovers therefore do not drive the baseline estimate.

[49] I use all spikes to define S_{mt}, not only firms' first spikes as in the main DiD design. The event study itself is still estimated around the first spike. The sector is defined at the two-digit level.

TABLE C6. Local spillovers: the share of spiking firms in the same municipality, sector, and year.

	(1) Baseline	(2) Baseline $+ S_{mt}$	(3) Not-yet-treated firms	(4) Not-yet-treated firms, lags
Treated	−1.64***	−1.82***		
× Post	(0.57)	(0.60)		
S_{mt}		−0.01	−0.05	−0.18
		(0.30)	(0.32)	(0.31)
$S_{m,t-1}$				−0.25
				(0.19)
$S_{m,t-2}$				−0.14
				(0.19)
$S_{m,t-3}$				−0.32
				(0.48)
Units	pp	pp	pp	pp
Fixed effects				
Firm-event-year	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Observations	263k	230k	200k	105k

Notes: The table reports ordinary least squares estimates of the pooled difference-in-differences specification in Equation (6) with controls for local spillovers. S_{mt} is the share of firms in the firm's municipality and sector m that have a DT investment spike, as defined in Equation (5), in year t. Column 1 repeats the baseline. Column 2 adds S_{mt} to the baseline regression. Columns 3 and 4 restrict the stacked sample to not-yet-treated firms and regress their early-career employment share on S_{mt} (Column 3) and on S_{mt} and its three lags (Column 4). The outcome in every column is the early-career employment share in percentage points (0–100), so coefficients on S measure the change in the share, in percentage points, when the local spike share rises from 0 to 1. The sample is the stacked panel described in Section 3.2, with event years s from 2010 to 2019; Columns 2 to 4 lose firm-years for which S_{mt} or its lags are not observed. Observations are weighted by the firm's early-career employment at $k = -1$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

C.8. Early-Career Employment at IT and Business-Services Firms

Investment spikes in DT could result in outsourcing of IT and back-office functions, which I do not observe in the data. Appendix C.3 finds that there is no evidence of an *increase* in outsourcing following an investment spike. Here, I extend the local spillover analysis from Appendix C.7 to show that early-career employment at IT and business-services firms—the providers of DT automation services—does not rise in response to spikes within their municipalities.[50] Let S_{mt} denote the share of firms in municipality m that have a spike in year t. Then I take all IT and business-services firms from the matched employer–employee data (not the stacked DiD) and regress their log early-career employment on S_{mt} and three of its lags, including firm and year fixed effects. Each firm is weighted by total employment. Table C7 shows that none of the coefficients are positive or significant, which implies there is no evidence that service providers absorb non-hired early-career workers following investment spikes. Note, the estimates are very noisy due to the small sample size.

[50] I no longer define the variable S within the same sector because I treat IT and business-services firms as the sink for non-hired early-career workers.

TABLE C7. Early-career employment at IT and business-services firms and local spike share.

	(1) ln(Early-career employment)
S_{mt}	−37.37*
	(22.50)
$S_{m,t-1}$	−5.36
	(18.58)
$S_{m,t-2}$	0.77
	(18.80)
$S_{m,t-3}$	−19.27
	(24.04)
Units	log points
Fixed effects	
Firm	Yes
Year	Yes
Observations	2,341
Within R-squared	0.003

Notes: The table reports ordinary least squares estimates from a regression of log early-career employment at IT and business-services firms on S_{mt}, the share of firms in the firm's municipality m with a digital-technologies investment spike, as defined in Equation (5), in year t, and on three lags of S_{mt}. The outcome is multiplied by 100 so that estimates are in log points per unit increase in the local spike share. The sample is firm-years of IT and business-services firms (SBI sectors 62–63, 69–70, 78) in the matched employer–employee data over the event years of the stacked design, 2010 to 2019. The regression includes firm and year fixed effects. Each firm is weighted by total employment. Standard errors, in parentheses, are clustered at the municipality level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

C.9. Estimates by Sector

Table C8 re-estimates the baseline DiD separately for firms in services, wholesale and retail trade, and industry, in levels and in logs. The EC share falls in every sector after a spike. The decline is largest in wholesale and retail (−2.5 pp, or 19 log points) and smallest in industry (−0.6 pp, or 6 log points), where the EC share is lower to begin with. The decline is therefore not confined to a single sector.

TABLE C8. Difference-in-differences estimates by sector.

	(1)	(2)	(3)	(4)	(5)	(6)
	Services		Wholesale and retail		Industry	
	Share	Log	Share	Log	Share	Log
Treated	−1.71*	−9.77**	−2.48***	−18.58***	−0.59**	−5.53
× Post	(0.88)	(4.33)	(0.71)	(5.56)	(0.30)	(4.40)
Units	pp	log points	pp	log points	pp	log points
Fixed effects						
Firm-event-year	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Observations	104k	97k	76k	69k	81k	73k

Notes: The table reports the pooled difference-in-differences estimate $\hat{\theta}$ from Equation (6), estimated separately by sector. Columns 1 and 2 use firms in services, Columns 3 and 4 firms in wholesale and retail trade, and Columns 5 and 6 firms in industry. The outcome is the early-career employment share in percentage points (0–100) in odd columns and the log of the early-career employment share, multiplied by 100 so that estimates are in log points, in even columns; the log specification drops firm-years with no early-career workers. Each sector re-stacks the panel of treated and not-yet-treated firms as described in Section 3.2, with event years s from 2010 to 2019. Observations are weighted by the firm's early-career employment at $k = -1$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

C.10. Worker-Level Digital Technologies Robustness

TABLE C9. Heterogeneity in the within-firm experience premium change.

	(1) Baseline	(2) Education	(3) Exposed field	(4) Concavity	(5) Hiring year
Premium: $\hat{\beta}$	0.139^{***}	0.135^{***}	0.136^{***}		0.116^{***}
	(0.006)	(0.006)	(0.006)		(0.008)
Post-event premium: $\hat{\gamma}$	0.023^{***}	0.001	0.016^{**}		0.038^{***}
	(0.007)	(0.007)	(0.007)		(0.010)
Post-event premium × group: $\hat{\gamma}^G$		0.057^{***}	0.024^{***}		
		(0.004)	(0.004)		
Premium, bin 2: $\hat{\beta}_2$				0.199^{***}	
				(0.011)	
Premium, bin 3: $\hat{\beta}_3$				0.282^{***}	
				(0.014)	
Post-event premium, bin 2: $\hat{\gamma}_2$				0.035^{***}	
				(0.012)	
Post-event premium, bin 3: $\hat{\gamma}_3$				0.041^{***}	
				(0.014)	
Fixed effects					
Firm-year	Yes	No	No	Yes	Yes
Firm-year × group	No	Yes	Yes	No	No
Demographic	Yes	No	No	Yes	Yes
Demographic × group	No	Yes	Yes	No	No
Observations	875k	847k	844k	875k	291k

Notes: The table reports ordinary least squares estimates of the pooled specification in Equation (9) and several extensions. The dependent variable is the log hourly wage. Column 1 repeats the baseline from Table 4. Columns 2 and 3 add the triple interaction $I_H \times I_P \times G_i$, in which G_i indicates a college degree (Column 2) or a field of study exposed to digital technologies (Column 3; Appendix A.4). Column 4 replaces the single indicator I_H with indicators for $\tau \in \{4, 5, 6\}$ (bin 2) and $\tau \in \{7, \ldots, 10\}$ (bin 3), relative to $\tau \in \{0, \ldots, 3\}$. Column 5 restricts the sample to each worker's hiring year at firm j. Demographic fixed effects are education level-by-field, gender, and migrant background. The sample is early-career workers from the 2006–12 cohorts, followed for $\tau \in \{0, \ldots, 10\}$, employed at treated or not-yet-treated firms (Section 3) within event time $k \in \{-4, \ldots, 3\}$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

TABLE C10. Robustness of the within-firm premium change to match-quality controls.

	Match-quality control	premium $\hat{\beta}$	Post-event premium $\hat{\gamma}$	Control coefficient	Fixed effects (added)	Obs.
(1)	Baseline	0.139^{***}	0.023^{***}			875k
		(0.006)	(0.007)			
(2)	Log hiring wage	0.076^{***}	0.012^{**}	0.457^{***}		874k
		(0.007)	(0.005)	(0.031)		
(3)	Cohort pct. of entry wage	0.139^{***}	0.023^{***}	0.244^{***}		835k
		(0.006)	(0.007)	(0.041)		
(4)	Cohort vigintile of entry wage	0.138^{***}	0.023^{***}		Vigintile	835k
		(0.006)	(0.007)			
(5)	Within-firm rank of entry wage	0.064^{***}	0.018^{***}	0.530^{***}		875k
		(0.007)	(0.005)	(0.042)		
(6)	Worker–firm match	-0.026^{***}	-0.028^{***}		Match	687k
		(0.003)	(0.003)			

Notes: The table reports ordinary least squares estimates of $\hat{\beta}$ and $\hat{\gamma}$ from the pooled specification in Equation (9) with controls for match quality measured at hiring. Each row is a separate regression. The dependent variable is the log hourly wage. Row 1 repeats the baseline from Table 4. Row 2 adds the log hourly wage in the year the worker is hired by firm j. Row 3 adds the percentile of the worker's wage at $\tau = 0$ within their cohort (continuous). Row 4 replaces the percentile with fixed effects for its vigintiles. Row 5 adds the rank of the hiring wage among workers hired by the same firm in the same year. The column 'Control coefficient' reports the estimate on the added continuous control. Row 6 adds worker-by-firm match fixed effects, which absorb the demographic fixed effects. Rows 1 to 5 include firm-by-year fixed effects and demographic fixed effects for education level-by-field, gender, and migrant background. Row 6 includes firm and year fixed effects separately. The sample is early-career workers from the 2006–12 cohorts, followed for $\tau \in \{0, \ldots, 10\}$, employed at treated or not-yet-treated firms (Section 3) within event time $k \in \{-4, \ldots, 3\}$. Standard errors, in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Appendix D. Quantitative Model

D.1. Equilibrium and Modelling Choices

Here I discuss dividends, the equilibrium definition, and the modelling choices I make for the baseline calibration.

D.1.1. Aggregate Dividends and Consumption

Each firm pays out dividends that equal operating profit net of investment:

$$d_{jt} = Y_{jt} - w_{Ejt}L_{Ejt} - w_{Hjt}L_{Hjt} - q_t x_{jt} \quad \Longrightarrow \quad D_t = \int d_{jt}dj. \tag{D1}$$

Substituting D_t into the household budget constraint in Equation (11) and using that each firm-type labour market clears $n_{ijt} = L_{ijt}$ in equilibrium:

$$C_t = \int Y_{jt}dj - q_t \int x_{jt}dj. \tag{D2}$$

D.1.2. Stationary Recursive Equilibrium

Let $P(\varepsilon, \varepsilon')$ denote the transition function of the effectiveness process in Equation (14), and let μ be a probability measure over firm states (X, ε). There are no wealth effects in labour supply, $\gamma = 0$, and I assume that $\kappa_i = 1$ for $i \in \{E, H\}$.

DEFINITION 1. *Given a collection of exogenous parameters and the investment price q, a stationary equilibrium consists of firm policy functions $L_E(X, \varepsilon)$, $L_H(X, \varepsilon)$, and $x(X, \varepsilon)$, wage (labour supply) functions $w_E(X, \varepsilon)$ and $w_H(X, \varepsilon)$, and a probability measure μ over firm states (X, ε) such that:*

1. *Firm optimisation: taking the labour supply functions and price q as given, the policy functions solve the firm problem in Equation (16).*
2. *Labour market clearing: each firm-type-specific labour market clears, $n_i(X, \varepsilon) = L_i(X, \varepsilon)$. The wage sits on both the demand and supply schedules: $w_i(X, \varepsilon) = MP_i(X, \varepsilon) = L_i(X, \varepsilon)^{1/\varphi_i}$ for $i \in \{E, H\}$, in which $MP_i \equiv \partial Y / \partial L_i$.*
3. *Stationarity: μ is invariant under the transition induced by the law of motion $X'(X, \varepsilon) = (1 - \delta)X + x(X, \varepsilon)$ and the transition function $P(\varepsilon, \varepsilon')$.*

D.1.3. Discussion of Modelling Choices

Here I compare various model features, especially the simplifying assumptions that I make, to the literature.

'Four Ss' of intangible capital. Haskel and Westlake (2018) identify four properties that separate intangible from tangible capital. The digital technologies that I cover in this paper are intangible technologies. The first 's' stands for *scalability*, or non-rivalry: an intangible asset can be used across several production streams at once. For example, the predictive software in the law firm can be used simultaneously by every senior (and junior) lawyer. In contrast, an industrial robot can only be used by a single worker at a time. In my model, effective DT M is partially non-rival because one additional unit (i) increases the productivity of *every* unit of post-early-career labour, while (ii) simultaneously substituting for early-career labour. Nevertheless, the first nest by itself features rivalry since it is homogeneous of degree one in (L_E, M).

The second 's' refers to *sunkenness*: reselling intangibles is difficult, especially because the value of an intangible asset is often tied up in the firm that deployed it. I impose sunkenness in the full irreversibility constraint, $x \geq 0$.

The third 's' represents *synergies*: the value of an intangible asset increases when combined with complementary assets, such as organisational capital and management practices. The effectiveness draw ε is an exogenous, reduced-form proxy for those complementary assets. Two firms with an identical stock X produce different levels of output because the effectiveness of their DT differs. Furthermore, DT are more valuable in production in firms with a higher employment of post-early-career labour through the augmentation channel.

The fourth property is *spillovers*. Haskel and Westlake (2018) argue that an investing firm cannot fully appropriate the value of an intangible asset, because (i) workers can carry some of the intangible-specific knowledge with them when they leave the firm, or (ii) competitors can imitate the asset. I abstract from spillovers in my model.

Productivity and scale. The effectiveness draw is the only exogenous source of cross-firm variation, so dispersion in output only arises from the current draw and the history of draws of ε. Other heterogeneous-firm models often include a Hicks-neutral productivity term to generate size dispersion (Hopenhayn and Rogerson 1993; Melitz 2003). The interaction of factor-neutral productivity shocks and irreversibility would also generate investment spikes following a sequence of favourable shocks. However, positive, yet neutral, shocks would raise the marginal product and labour demand of both labour types, increasing early-career employment after an investment spike, which is refuted by the data.[51]

D.2. Untargeted Moments

The calibration in Section 5 targets six moments. The estimated model also reproduces, to varying extents, several moments that I do not target.

[51] See Appendix C.4 for the results.

The 'log EC employment DiD' moment represents the DiD estimate from a regression in which I replace the EC employment share with the log EC employment. The results for this regression are shown in Column 2 of Table C2. The model shows a negative coefficient, just over a third as large as the untargeted moment from the data. The model also reproduces the positive 'log post-EC employment DiD' (albeit smaller) for the positive, yet insignificant, employment effect for non-EC workers, shown in Column 3 of Table C2.

The 'log EC wage DiD' moment represents the DiD estimates shown in Table 4. The results show how the early-career and post-early-career wage changes around investment spikes. The model does well to capture the decline in early-career wages and the increase in post-early-career wages. Note, in the data these two moments are calculated at the worker level without stacking; in the model they are calculated at the firm level with stacking.

The 'Share of event years' moment measures the share of observations in the 2006–12 firm panel with a DT investment event. The model underestimates the share by around 7 percentage points.

The 'Share of investment in event years' moment denotes the share of total DT investment from calendar year 2006 until 2012 that occurs in the event years. The model slightly overestimates the share compared to the data.

The 'Autocorrelation of investment / cost' moment represents the autocorrelation from a regression of the investment cost share on its lagged value.[52] The investment cost share is calculated as the ratio of investment expenditure to total operating costs. In the model, total operating costs are the sum of labour costs and investment costs. The model cannot reproduce the positive autocorrelation in the data, because the AR(1) effectiveness shock is mean-reverting.

The 'Labour share' moment is calculated as the sum of labour payments divided by the sum of value added in the data. I take the value from Silvestrini and Weder (2026) for the Dutch economy in the mid-2000s. In the model, the labour share is calculated as the sum of labour payments divided by the sum of output, which equals value added since there are no intermediate inputs.

[52] The regression includes sector-by-year fixed effects. In the simulated panel from the model, only year effects are included because there is no sectoral heterogeneity.

TABLE D1. Untargeted moments after calibration.

Moment	Source	Data	Model
Log EC employment DiD (log points)	Column 2 in Table C2	-11.68	-4.38
Log post-EC employment DiD (log points)	Column 3 in Table C2	2.38	1.35
Log EC wage DiD (log points)	Column 2 in Table 4	-1.54	-2.15
Log post-EC wage DiD (log points)	Column 3 in Table 4	0.42	0.68
Share of event years (share, 0-1)	2006–12 firm-year panel	0.12	0.05
Share of investment in event years (share, 0-1)	2006–12 firm-year panel	0.14	0.15
Autocorrelation of investment / cost	Regression described in text	0.56	-0.21
Labour share	Silvestrini and Weder (2026)	0.68	0.61

Notes: This table shows untargeted empirical moments from the data and the simulated model data after calibration. The two employment DiDs are taken from Table C2, in which I replace the EC employment share with the EC employment or the non-EC (post-EC) employment and re-run the DiD design. The two wage DiDs are taken from Table 4, which show EC and post-EC wage changes around investment spikes. In the data these are calculated at the worker level without stacking; in the model these are calculated at the firm level with stacking. The 'Share of event years' moment measures the share of observations in the 2006–12 firm panel with a DT event. The 'Share of investment in event years' measures the share of total DT investment from calendar year 2006 until 2012 that occurs in the event years. The 'Autocorrelation of investment / cost' moment represents the autocorrelation when I regress the investment-to-total-cost ratio on the lagged value, including sector-year fixed effects. 'Labour share' is calculated as the sum of labour payments divided by the sum of value added. I take the value from Silvestrini and Weder (2026).

D.3. Derivation of Elasticity of Substitution

I target the employment-weighted elasticity of substitution between the labour types, Ξ. The firm-level elasticity of substitution between the two labour types is given by:

$$e_{jt} = \frac{1}{1-\sigma}\left[\xi_{jt} + \left(1-\xi_{jt}\right)\left(1+\frac{\vartheta M_{jt}}{L_{Ejt}}\right)\right], \qquad \xi_{jt} \equiv \frac{\left(L_{Ejt}+\vartheta M_{jt}\right)^{\sigma}}{\left(L_{Ejt}+\vartheta M_{jt}\right)^{\sigma} + \left(M_{jt}^{\eta} L_{Hjt}\right)^{\sigma}}, \tag{D3}$$

$$\Xi = \frac{\sum_{j,t}\left(L_{Ejt}+L_{Hjt}\right)e_{jt}}{\sum_{j,t}\left(L_{Ejt}+L_{Hjt}\right)}. \tag{D4}$$

D.4. Derivation of Scale and Substitution Terms

The scale–substitution decomposition of $\ln L_E$ follows directly from the equilibrium in the early-career labour market. Every firm hires until the marginal product equals the wage on the

labour supply curve:

$$\nu A^{\nu/\sigma-1}\left(L_E+\vartheta M\right)^{\sigma-1}=L_E^{1/\varphi_E}, \tag{D5}$$

where $A\equiv\left(L_E+\vartheta M\right)^{\sigma}+\left(M^{\eta}L_H\right)^{\sigma}$ is the CES aggregator. Taking logs and collecting terms gives

$$\ln L_E=\frac{1}{1/\varphi_E-(\sigma-1)}\Big[\underbrace{\ln\left(\nu A^{\nu/\sigma-1}\right)}_{\text{scale}}+\underbrace{(\sigma-1)\ln\left(1+\frac{\vartheta M}{L_E}\right)}_{\text{substitution}}\Big]. \tag{D6}$$

D.5. Parameter Robustness

In this section, I report the counterfactual results for various robustness checks regarding externally set parameters.

Robustness to labour supply elasticities. The baseline calibration assumes the firm-level labour supply elasticities $\varphi_E=\varphi_H=2$. However, there are no estimates of firm-level elasticities that differ by experience (or age) in the literature. To proceed, I take three geometric means of φ_E and φ_H: $\bar{\varphi}\equiv\sqrt{\varphi_E\times\varphi_H}\in\{1,2,3\}$ and three ratios of $r\equiv\varphi_E/\varphi_H\in\{0.5,1,2\}$, which results in a 3×3 grid of elasticities. Panel A of Table D2 compares the counterfactual results from the baseline calibration to the other eight elasticity combinations.

TABLE D2. Counterfactuals under alternative parameters and model extensions.

Extension	Δ Returns to early-career experience (log points)	ΔL_E (%)
Panel A: Labour supply elasticities		
$\varphi_E = 0.71$, $\varphi_H = 1.41$	3.53	-1.69
$\varphi_E = 1.00$, $\varphi_H = 1.00$	3.89	-2.06
$\varphi_E = 1.41$, $\varphi_H = 0.71$	3.40	-3.16
$\varphi_E = 1.41$, $\varphi_H = 2.83$	2.89	-3.19
Baseline ($\varphi_E = 2.00$, $\varphi_H = 2.00$)	3.03	-4.51
$\varphi_E = 2.83$, $\varphi_H = 1.41$	3.36	-6.01
$\varphi_E = 2.12$, $\varphi_H = 4.24$	3.09	-4.84
$\varphi_E = 3.00$, $\varphi_H = 3.00$	3.04	-6.06
$\varphi_E = 4.24$, $\varphi_H = 2.12$	1.92	-2.34
Panel B: Depreciation rate and discount factor		
Baseline ($\delta = 0.20$, $\beta = 0.99$)	3.03	-4.51
$\delta = 0.10$	3.63	-4.56
$\delta = 0.40$	2.79	-4.17
$\beta = 0.96$	2.98	-3.58
Panel C: Model extensions		
Baseline	3.03	-4.51
CES	0.01	0.05
Wealth effects in labour supply ($\gamma = 2$)	3.13	-5.28
Fixed adjustment costs (Appendix **??**)	3.34	-4.84
Experience accumulation, long run (Appendix **??**)	3.37	-5.64

Notes: This table reports the counterfactual changes across steady states following the decline in the aggregate DT investment price. The first column reports the change in the returns to early-career experience in log points and the second column the % change in early-career labour L_E. Panel A varies the firm-level labour supply elasticities over a 3×3 grid of geometric means $\bar{\varphi} \equiv \sqrt{\varphi_E \times \varphi_H} \in \{1, 2, 3\}$ and ratios $r \equiv \varphi_E/\varphi_H \in \{0.5, 1, 2\}$. Panel B varies the depreciation rate δ and the discount factor β. Panel C reports the model extensions discussed in Appendix **??**. Each extension is re-estimated with the same SMM procedure.

Robustness to depreciation rate and discount factor. The baseline calibration assumes a depreciation rate of $\delta = 0.2$, which lies in the range from Peters and Taylor (2017). For example, the depreciation rates used in Peters and Taylor (2017) vary from around 10% in pharmaceuticals up to 40% for computers. The '$\delta = 0.10$' and '$\delta = 0.40$' rows in Panel B of Table D2 show the baseline result is robust if I set δ equal to the bounds.

One model period coincides with a year. In the baseline calibration, I let $\beta = 0.99$ to match the real interest rate of around 1%. Panel B of Table D2 shows that setting β lower, e.g., $\beta = 0.96$, changes the baseline result little.

D.6. Model Extensions

In this section, I discuss various model extensions and the robustness of the baseline counterfactual results.

Fixed Adjustment Costs. The baseline model generates lumpy investment from the interaction of the effectiveness draws ε_{jt} and irreversibility. In this extension, I replace the effectiveness draws with a random fixed cost of investing F_{jt}. The cost of investing x_{jt} then becomes:

$$c(x_{jt}, F_{jt}, X_{jt}) = q\,x_{jt} + F_{jt}\,\pi(X_{jt})\,\mathbb{1}\{x_{jt} > 0\}, \tag{D7}$$

in which $\pi(X_{jt})$ denotes operating profit. F_{jt} is an idiosyncratic cost shifter that enters the investment cost and follows an AR(1) process. In logs, the process has a persistent and a transitory part:

$$\ln F_{jt} = \mu_F + a_{jt} + s\,z_{jt}, \qquad a_{jt} = \rho_F a_{jt-1} + e_{jt}, \tag{D8}$$

in which z_{jt} is an i.i.d. standard normal and $e_{jt} \sim \mathcal{N}\left(0, \sigma_F^2\right)$. The persistent component a_{jt} is a state variable and is discretised as ε_{jt} is in the baseline. I include the transitory component z_{jt} to smooth the investment hazard, fixing $s = 0.5$ exogenously. At the estimates, the transitory component accounts for a small share of the dispersion in $\ln F_{jt}$. Given the AR(1) cost process, a firm invests whenever its draw F_{jt} falls below a threshold $F^*(X_{jt}, a_{jt})$, so the timing of an investment event depends only on the cost draw and the pre-event state.

I estimate $(\mu_F, \sigma_F, \rho_F)$ in place of $(\rho_\varepsilon, \sigma_\varepsilon^2)$ and add the investment event frequency to the model targets. The estimates are $\hat{\mu}_F = -6.47$, $\hat{\sigma}_F = 2.38$, $\hat{\rho}_F = 0.97$. Across steady states, the returns to early-career experience rise by 3.3 log points, whereas early-career labour falls by 4.8%, close to the baseline's 3.0 log points and 4.5%.

Experience Accumulation. Labour types are fixed in the baseline model. Here the stock of post-early-career workers accumulates as previously employed early-career workers gain

experience:

$$N_{Ht+1} = (1 - \delta_H)N_{Ht} + \lambda N_{Et}, \tag{D9}$$

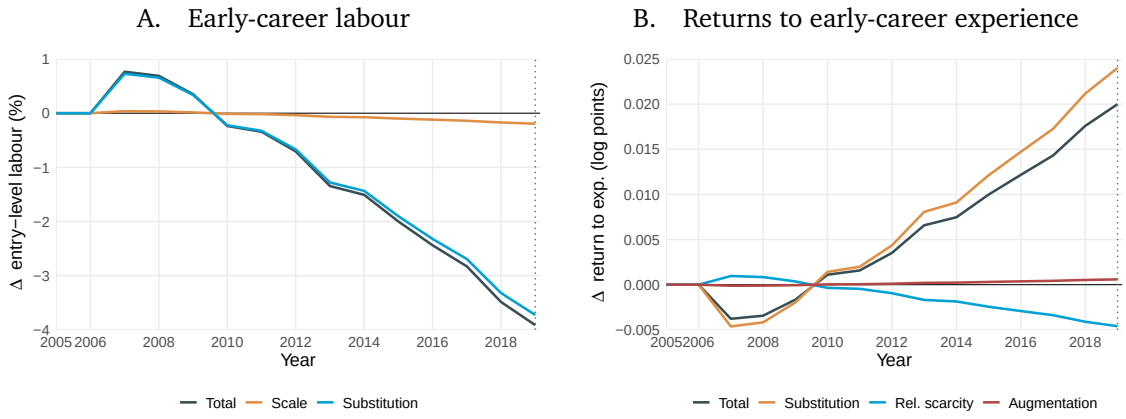
in which N_{Ht} is the stock of post-early-career workers available at time t, and N_{Et} is the number of *employed* early-career workers at time t. I set $\delta_H = 1/40$ for a 40-year working life and $\lambda = 1/5$, since an early-career worker becomes a post-early-career worker after five years. I assume that the stock of early-career labour is elastic as in the baseline. However, the post-early-career stock is predetermined each period, and shrinks when early-career employment falls.

The labour market for post-early-career labour needs to clear each period, so I include a common shifter κ_{Ht} on the post-early-career labour supply curve, $w_{Ht} = \kappa_{Ht} L_{Ht}^{1/\varphi_H}$, which adjusts each period to clear the market. I re-estimate the extension with the same SMM procedure and target an economy-wide post-early-career to early-career return of 0.35, which I estimate as the average within-person wage growth from $I_H = 0 \to I_H = 1$, where $I_H = 1$ if $\tau_{it} > 5$ and zero otherwise.

I implement a gradual price decline from 2006 to 2019 as the counterfactual, which splits the total 3.6% price decline evenly across each year. By 2019, early-career labour declines by 3.91% (baseline across steady states: 4.51%) and the experience premium rises by 2.00 log points (baseline: 3.03).

Even though the price decline is gradual and monotone, the firms internalise the price decline and postpone investment initially. Thus, early-career employment rises by 0.8% in the first years and only reaches its long-run level once the price stops falling; the long-run outcomes are the same.

FIGURE D1. Dynamic effects of a gradual price decline.



Notes: Dynamic effects of a gradual price decline in the extension that allows for experience accumulation of workers.

Wealth effects in labour supply. Setting $\gamma = 2$ in the household's utility function introduces wealth effects in labour supply. The returns to early-career experience rise by 3.1 log points against 3.0 log points in the baseline calibration. The employment decline is deeper, 5.3% against 4.5%, because the accumulation of DT raises household income and the wealth effect lowers labour supply on top of the demand shift.

Constant-elasticity-of-substitution (CES). The baseline imposes the substitution–augmentation asymmetry on the production function in Equation (13), guided by economic intuition and the DiD evidence. Following Krusell et al. (2000) and Jaimovich et al. (2013), I leave both elasticities unrestricted and discipline them with the data. The production function becomes:

$$Y_{jt} = \left(\mu L_E^{\rho_E} + (1-\mu)\left(\lambda M^{\rho_H} + (1-\lambda) L_H^{\rho_H}\right)^{\frac{\rho_E}{\rho_H}} \right)^{\nu/\rho_E}, \quad \text{with } -1 < \rho_E, \rho_H < 1 \text{ and } 0 \le \mu, \lambda \le 1, \tag{D10}$$

with *jt* subscripts dropped to simplify notation. I estimate $\hat{\rho}_E = 0.84$, $\hat{\rho}_H = 0.49$, $\hat{\mu} = 0.19$, $\hat{\lambda} = 0.58$. The estimates $\hat{\rho}_E, \hat{\rho}_H > 0$ imply both labour types are gross substitutes with DT. However, the unrestricted model fails an untargeted moment that the baseline matches better: after an investment event, EC employment falls by 11.7% in the data, 4.4% in the baseline model. In the CES extension, EC employment *rises* by 0.002%. Since both labour types are relatively substitutable with DT, the marginal products and wages decline enough to offset the substitution effect. Across steady states the unrestricted model delivers a 0.01 log point rise in the return and a 0.05% *increase* in early-career labour.